MAY/JUNE 2025 • LEVEL UP

THE MACHINE IS THE MESSAGE

WIRED

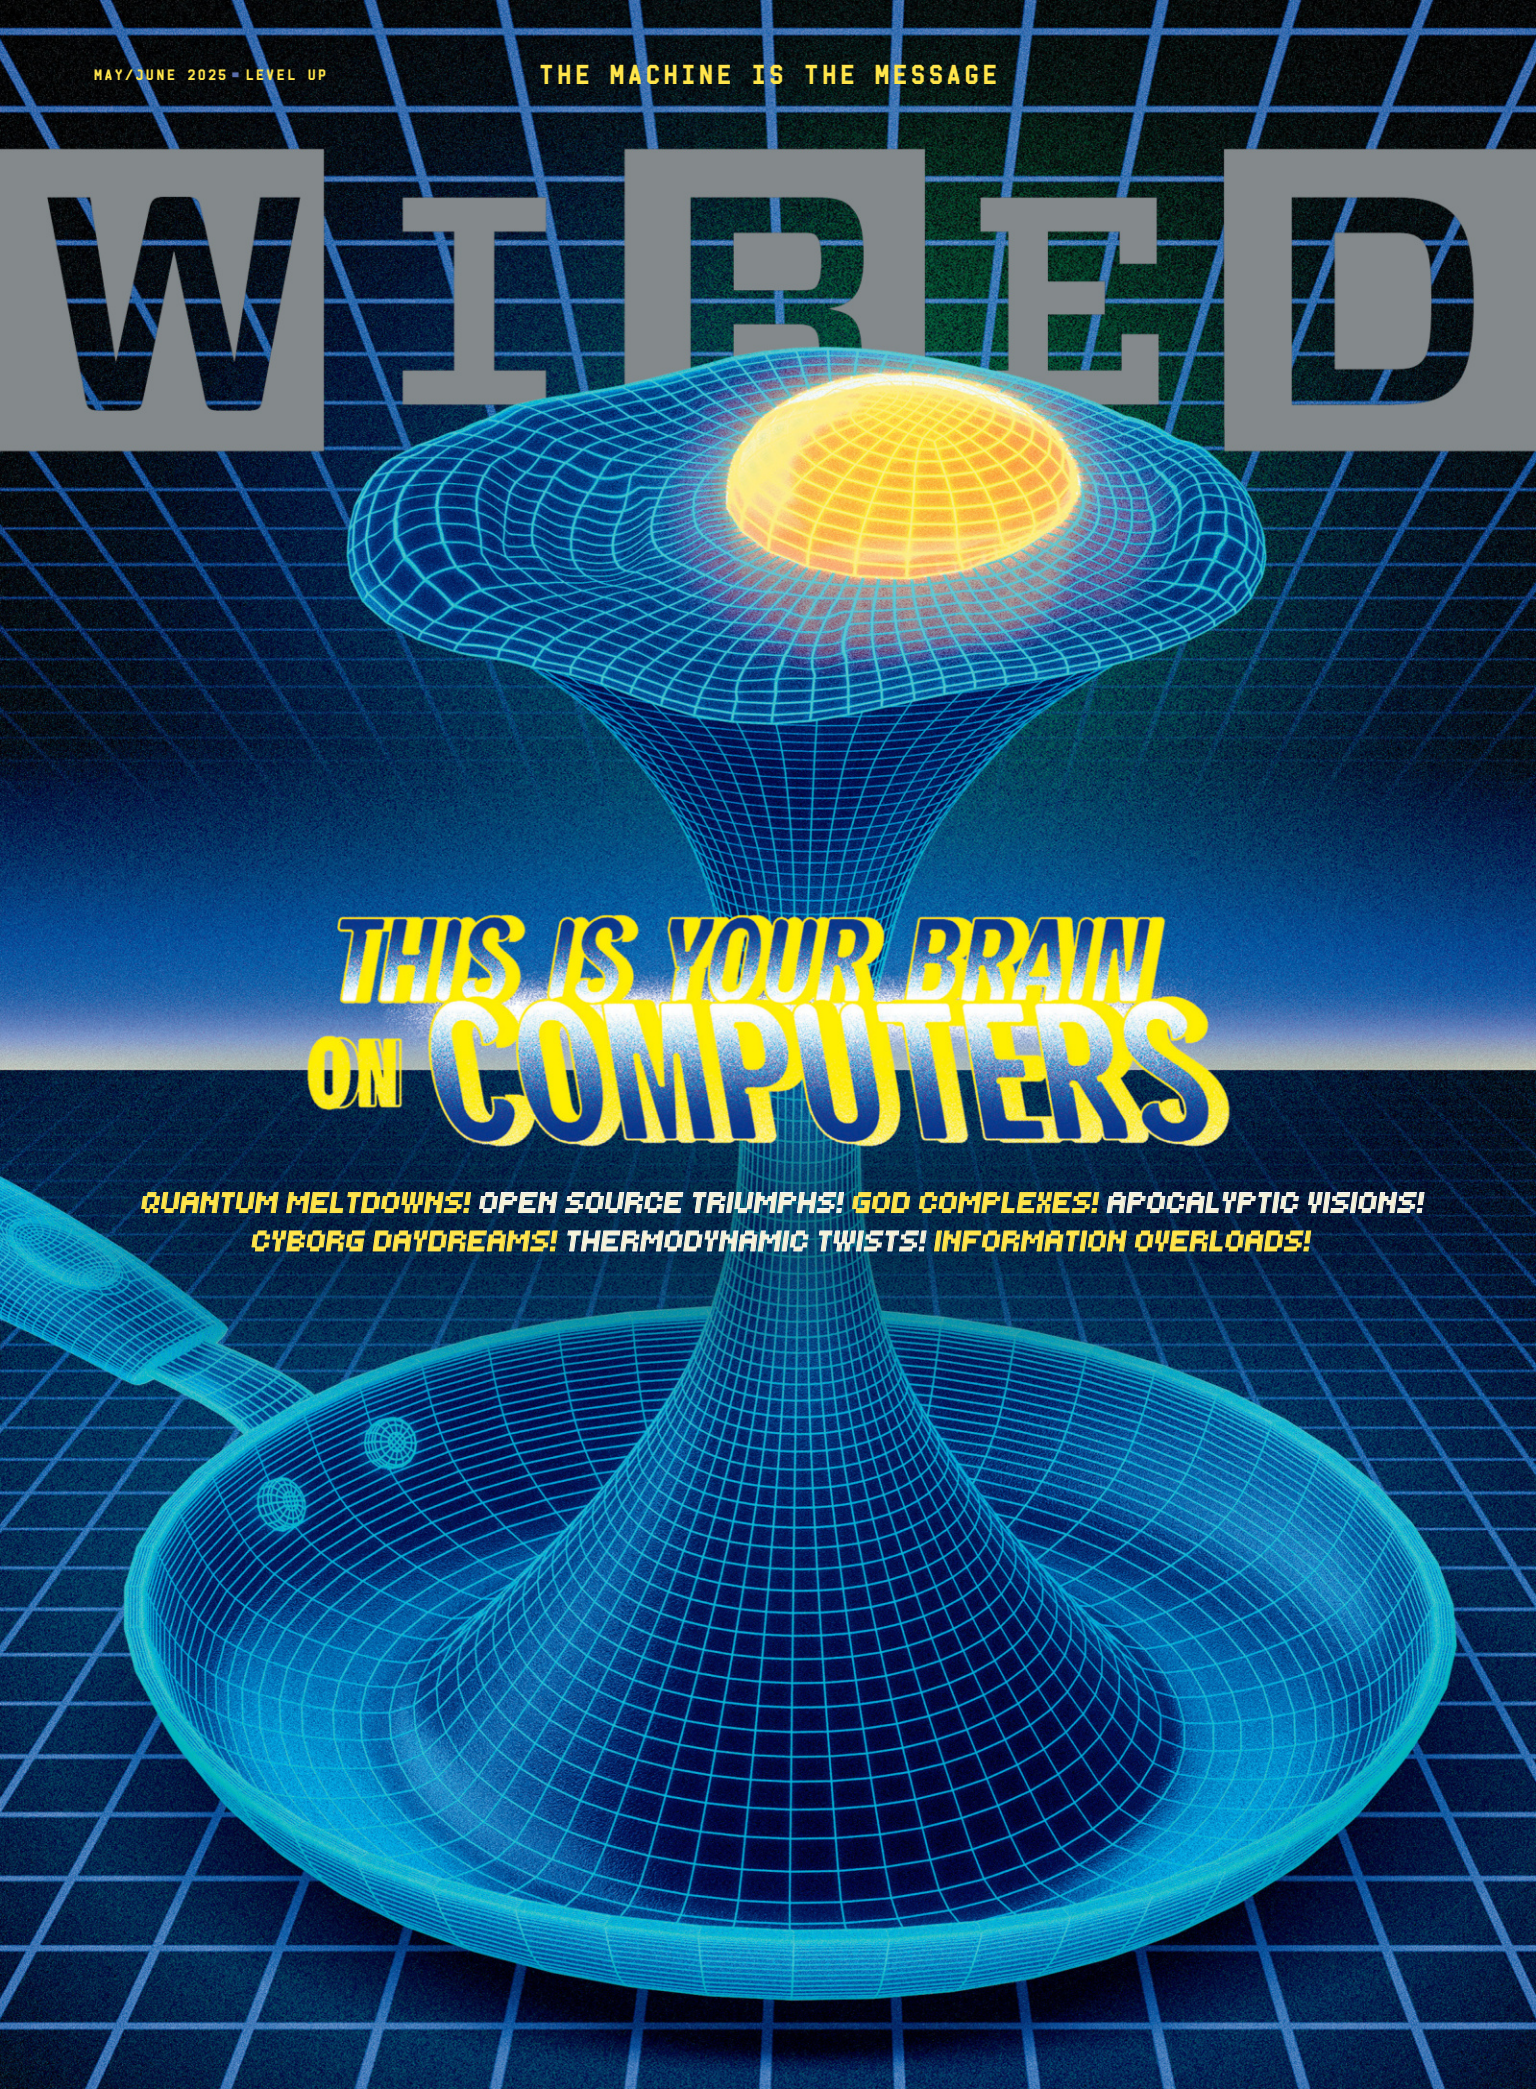

THIS IS YOUR BRAIN ON COMPUTERS

QUANTUM MELTDOWNS! OPEN SOURCE TRIUMPHS! GOD COMPLEXES! APOCALYPTIC VISIONS! CYBORG DAYDREAMS! THERMODYNAMIC TWISTS! INFORMATION OVERLOADS!

OMEGA
CO-AXIAL
MASTER CHRONOMETER
300m / 1000ft

omegawatches.com

PRECISION. POWER. POISE.

Set your eyes on the new Seamaster Diver 300M, with its monochromatic steel and Master Chronometer precision. A fusion of metals forged for endurance. Wear it and wear a legacy worn on some of the most daring wrists. Your next adventure won't go unnoticed. Dive deeper, go further, and let the Seamaster lead the way.

OMEGA

Readers weep for the wonder drug pipeline and ponder the solar sales abyss.

In our March/April issue, Bradley Hope profiled the United Arab Emirates' spy chief, Sheikh Tahnoun bin Zayed al Nahyan, and his attempts to become the AI industry's indispensable man. Virginia Heffernan chronicled Novo Nordisk's agonizing attempts to balance demand for its blockbuster products Ozempic and Wegovy—treated as weight-loss miracle drugs in the US—against its lifesaving treatments for type 1 diabetes. And Brendan I. Koerner followed one college student's journey into the alluring, all-devouring world of door-to-door solar panel sales.

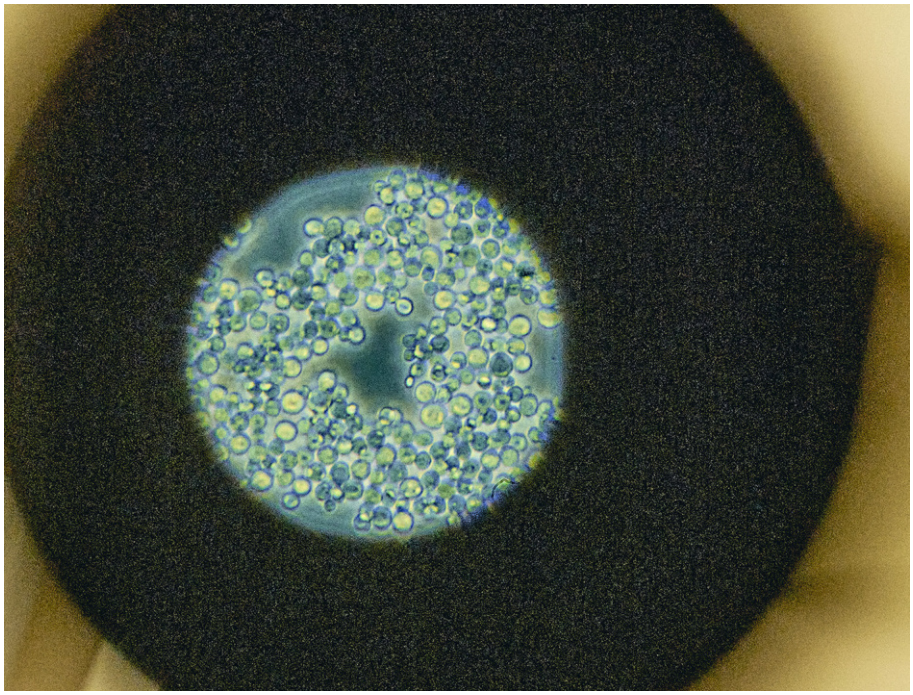

STEVEN ACHIAM

RE: "THE SHEIKH'S GAMBIT"

We tend to think in American-centric ways when considering the power of the Tech Bro community—this is a mistake.
—@ascablander.bsky.social

Boy, as Americans, we are screwed! Appreciate the research and the way it was presented, but it is now another danger/crisis on the horizon—no—at our doorstep.
—Deborah Thompson

RE: "THE KING OF OZEMPIC IS SCARED AS HELL"

Can't remember the last time a *news* article brought me to tears only because it was so well written, with a humbleness and sincerity and depth that we thirst for in the fast-paced world today.
—Aja Humm

I am very troubled by this article and for the health workers who are so diligently supporting the struggle of people with type 1 diabetes. That one market can have that much of a hold over the rest of the world just shows us how vulnerable we really are. I do hope the CEO of Novo can see a way through that will not disadvantage wonderful humans like my sister, who need the insulin to live.
—Tracey Winter

I understand how they are caught between a rock and a hard place. Everybody wants to be on the winning side financially, and that inevitably leads to loss in unexpected quarters.
—Elaine Saksons

RE: "A PLACE IN THE SUN"

This should be required reading served before every algorithmic promotion of hustler culture influencer content.
—@wilkos.bsky.social

Interesting piece about the "green economy's forward trenches" and the exploitation of bro-salesmen. We often hear about the virtues of the energy transition but rarely about the manpower propelling it.
—@barratt_dewey

This specific brand of sales bro is all over the place right now. My nephew is 24, and if he wasn't highly skeptical and a little introverted, I'd be more worried about him.
—u/running_hoagie

GET MORE WIRED
All WIRED stories can be found online, but only subscribers get unlimited access. If you are already a print subscriber, you can authenticate your account at WIRED.com/register.

RE: "THE KING OF OZEMPIC IS SCARED AS HELL"

"I do not envy the CEO of Novo."

—u/you_were_mythtaken

BUILT TO
PROTECT
TO EQUIP FOR
THE MISSION

CREATORS
RESPONDERS
DEFENDERS
EXPLORERS
TRAVELERS
PROTECTORS
PRODUCERS
ADVENTURERS
HUNTERS
GATHERERS
CAMPERS

TOP VIEW

ATX

INTRODUCING PELICAN ATX™ – ADVENTURE TRAVEL X-OVER

Legendary Pelican™ protection designed for the adventurous traveler who demands luggage that provides ultimate protection and built-in organization.

// Pelican Armor™ // Pelican Organization // Pelican TSA Latches // Pelican Mission Approved

©2025 Pelican Products, Inc.

PELICAN™ | TRVL

33.05

THE MACHINE IS THE MESSAGE

STORIES FROM THE FRONTIERS OF COMPUTING

Tech companies don't care whether you know how their products work. In fact, they don't *want* you to know. The more you assume technology is conjured into existence by a priest class of techno-mystics whose powers dwarf your own, the more you allow yourself to be controlled by it (and them). Don't let this happen. Though they seem to get more brain-fryingly complicated every year, computers are not, ultimately, magical devices. They are physics plus engineering, and they're easy enough to break down. In this special issue—which we've arranged along a computer's famous "layers of abstraction," from the bitwise logic of circuits and gates up to the fancy products themselves—we go deep inside the modern machine, where incredible things are happening. Freaky-deaky things, yes, but grokkable just the same. So take the plunge. It's time to get computers, people—before computers get you. –THE EDITORS

OPER
SYST

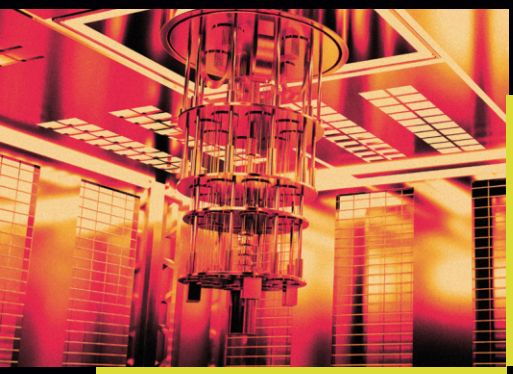

P.18
ALL HAIL THE GOD OF THERMO HARDWARE
↓ by Will Knight

ARCHITECTURES

P.8
THE QUANTUM HOLIDAY TO END ALL HOLIDAYS
↓ by Amit Katwala

GATES + CIRCUITS

P.26
ANGELINA JOLIE WAS RIGHT:
RISC CHANGED EVERYTHING
↑ by Jason Kehe

P.16
SCHRÖDINGER'S COMPUTER
← by WIRED staff

PHYSICS

P.24
HOW MUCH DOES THE INTERNET WEIGH?
↖ by Samantha Spengler

P.64
INSIDE ANTHROPIC
→ by Steven Levy

PRODUCTS

P.52
THE IMPOSSIBLE QUEST TO SIMULATE A WORM
↓→ by Claire L. Evans

ALGORITHMS + APPLICATIONS

P.56
GEAR THAT BOOSTS YOUR HUMANITY
↑ by the WIRED Gear team

PROGRAMMING LANGUAGES

P.40
HOW CODERS REALLY USE AI
←↑ by WIRED staff

P.44
MEET THE ARXIVIST
↖↑ by Sheon Han

ATING
EMS

P.34
PROGRAMMING AT THE END OF THE WORLD
←↑ by Tiffany Ng

ON THE COVER

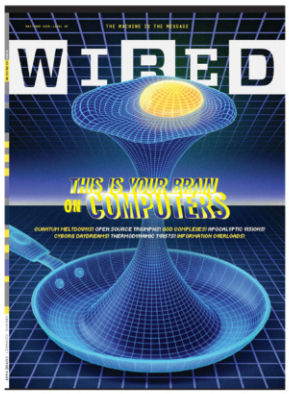

Drugs fry your brain like an egg, remember? We had one of our favorite illustrators, **Alvaro Dominguez**, reimagine that iconic bit of ’80s propaganda for the digital age. Are computers the new drugs? After you’re done with this issue, tell us how your brain feels.

DAN WINTERS; NICHOLAS LAW; JOHN PROVENCHER

P.74
MAX IMPACT:
THE WORST 7 YEARS IN BOEING’S HISTORY
by Lauren Smiley

P.88
SIX-WORD SCI-FI
by WIRED readers

The future of cool, comfortable sleep is here

Sleep up to

15°

Cooler on each side†

Meet the ClimateCool™ smart bed.

The only smart bed in the world that actively cools and effortlessly adjusts to both of you.*
It cools up to 20x faster than a leading competitor.† So you sleep comfortably all night.

sleep number®

visit sleepnumber.com/wired | 1-877-316-3922

*Based on use of temperature features and our proprietary Responsive Air® technology. †Based on internal microclimate testing.

ART BY VALENTINA CASALI

PHOTOGRAPH BY MITCH PAYNE

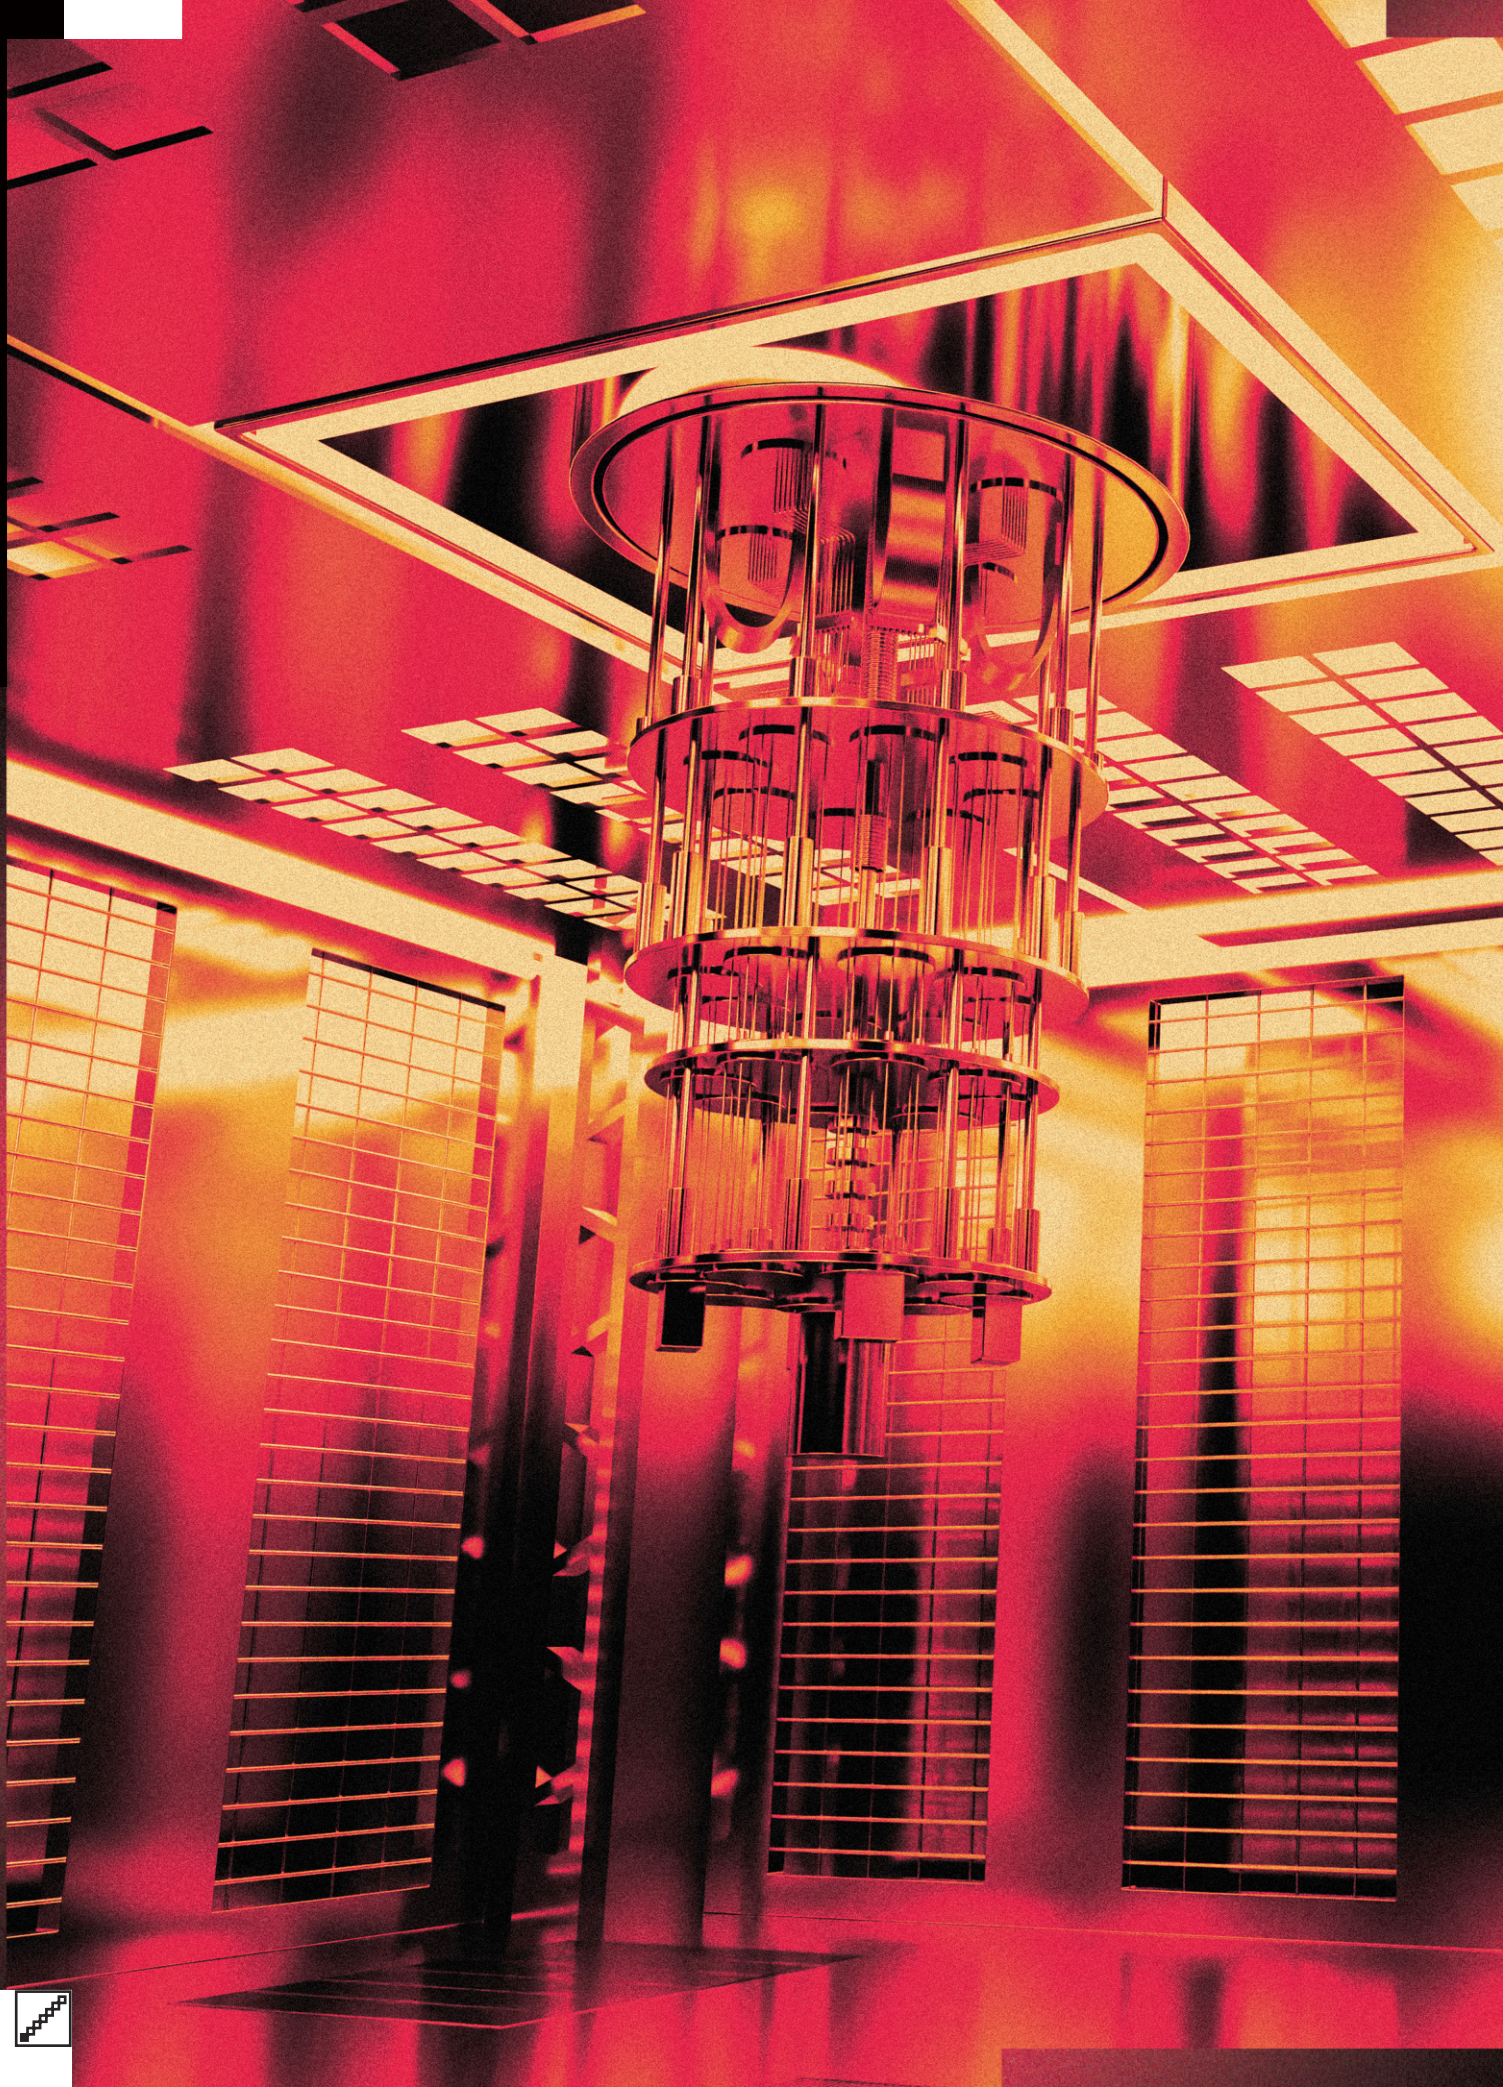

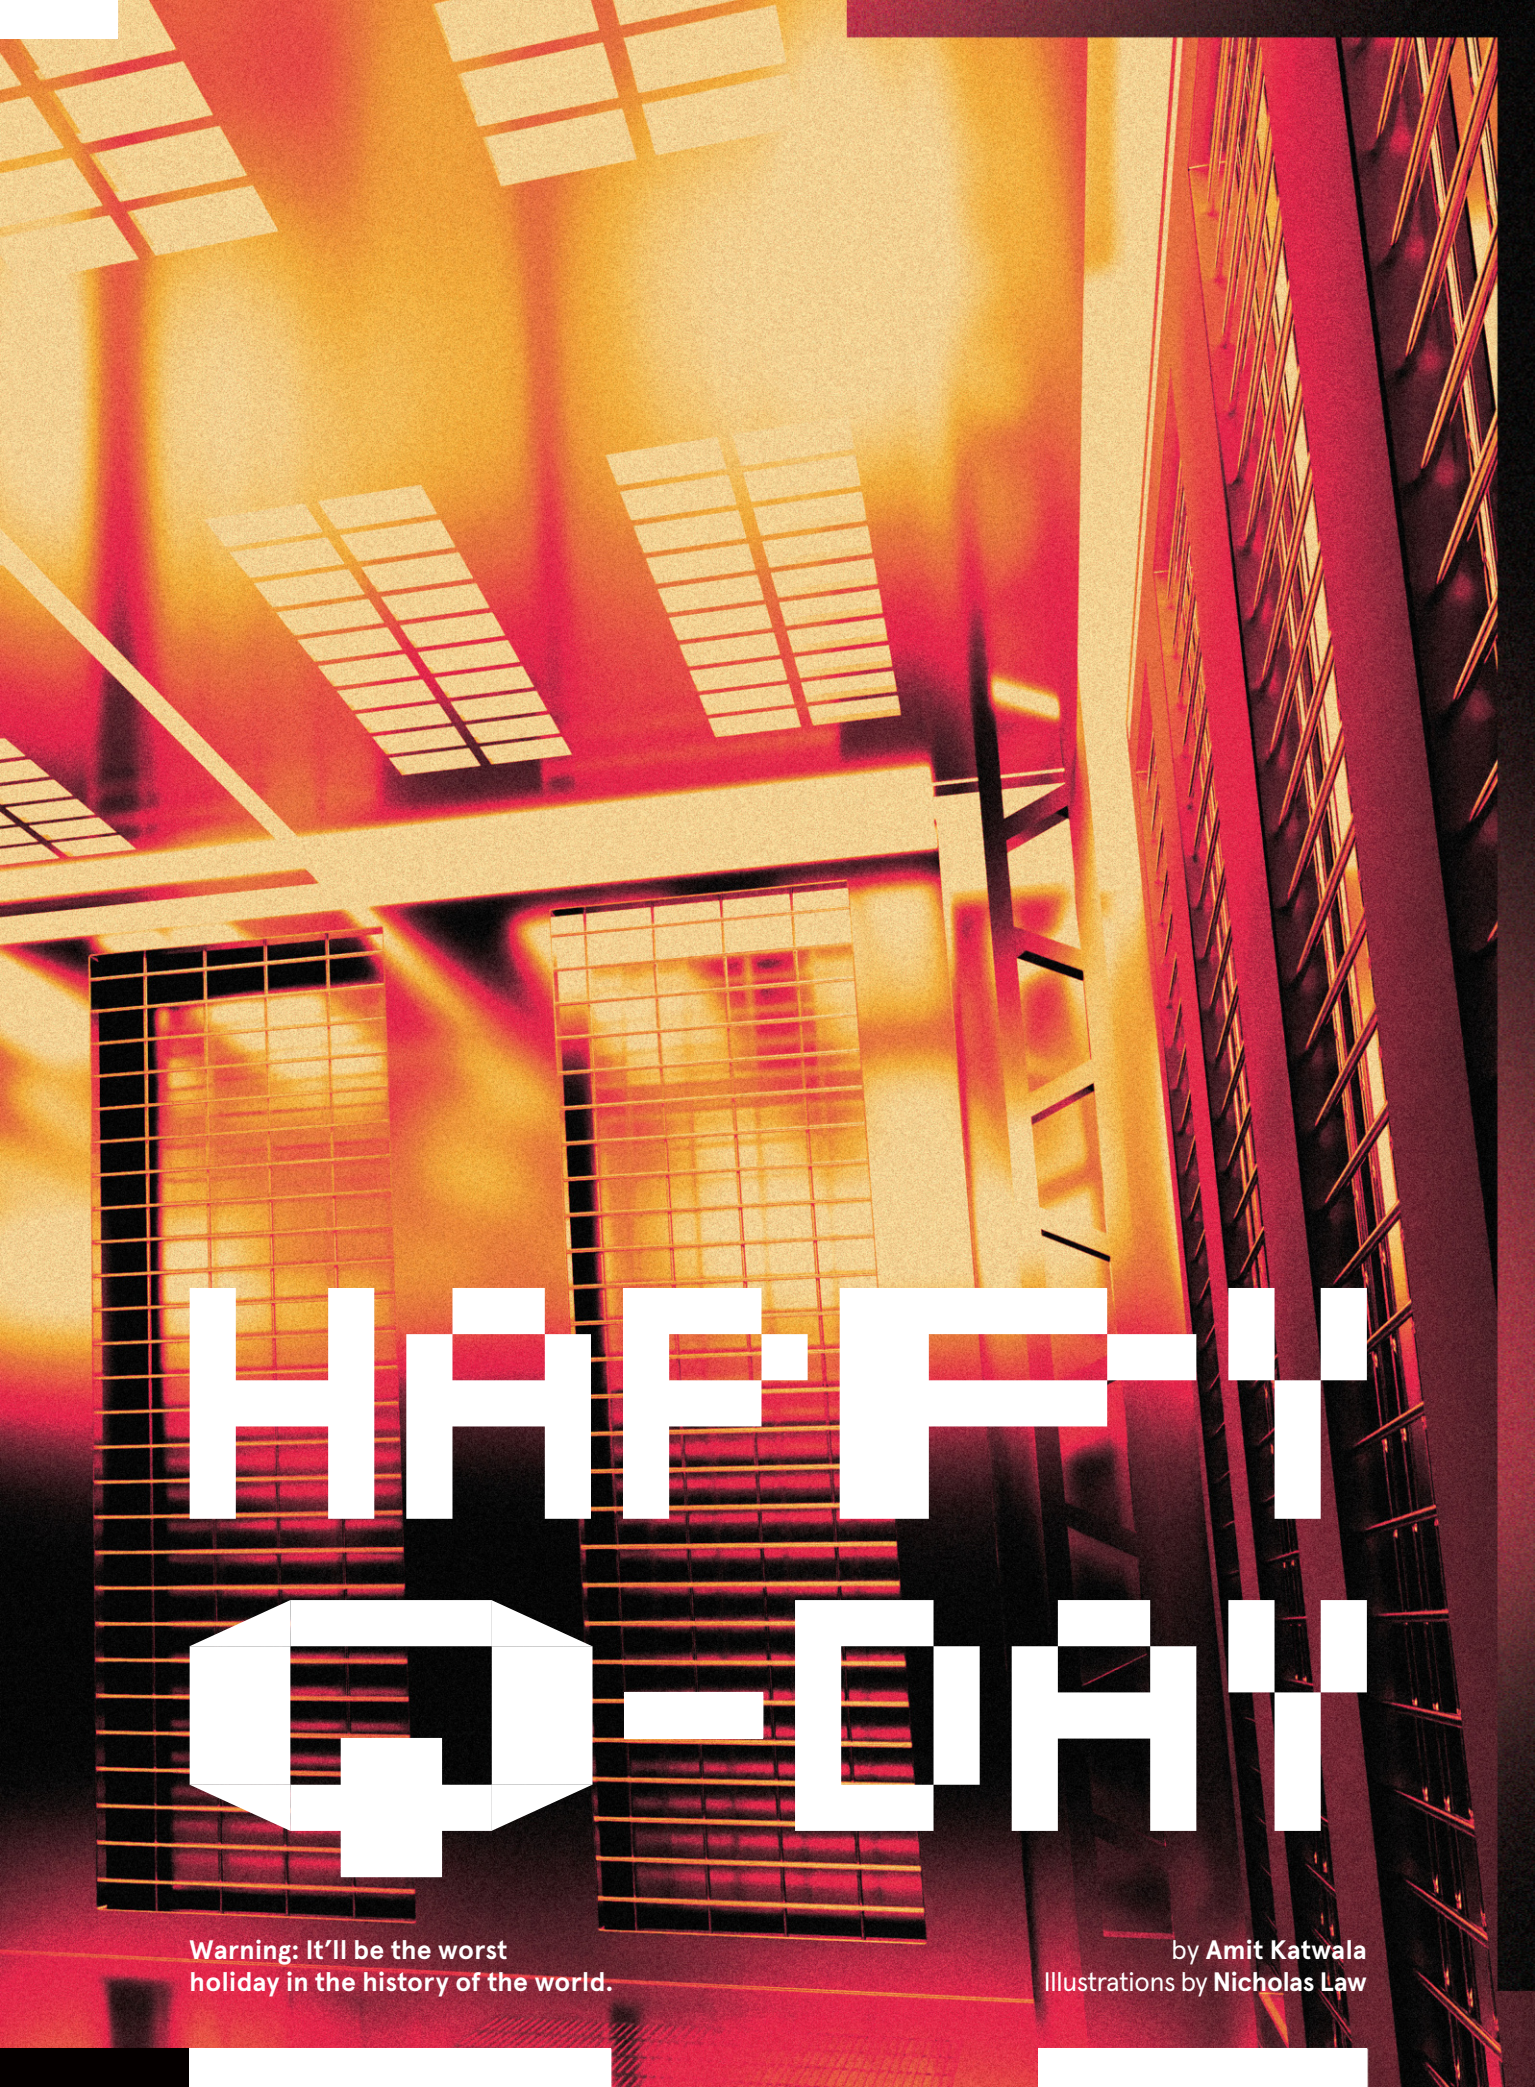

HALF-A-DAY

Warning: It'll be the worst holiday in the history of the world.

by **Amit Katwala**
Illustrations by **Nicholas Law**

One day soon, at a research lab near Santa Barbara or Seattle or a secret facility in the Chinese mountains, it will begin: the sudden unlocking of the world's secrets. Your secrets.

Cybersecurity analysts call this Q-Day—the day someone builds a quantum computer that can crack the most widely used forms of encryption. These math problems have kept humanity's intimate data safe for decades, but on Q-Day, everything could become vulnerable, for everyone: emails, text messages, anonymous posts, location histories, bitcoin wallets, police reports, hospital records, power stations, the entire global financial system.

"We're kind of playing Russian roulette," says Michele Mosca, who coauthored the most recent "Quantum Threat Timeline" report from the Global Risk Institute, which estimates how long we have left. "You'll *probably* win if you only play once, but it's not a good game to play." When Mosca and his colleagues surveyed cybersecurity experts last year, the forecast was sobering: a one-in-three chance that Q-Day happens before 2035. And the chances it has *already* happened in secret? Some people I spoke to estimated 15 percent—about the same as you'd get from one spin of the revolver cylinder.

The corporate AI wars may have stolen headlines in recent years, but the quantum arms race has been heating up too. Where today's AI pushes the limits of classical computing—the kind that runs on 0s and 1s—quantum technology represents an altogether different form of computing. By harnessing the spooky mechanics of the subatomic world, it can run on 0s, 1s, or anything in between. This makes quantum computers pretty terrible at, say, storing data but potentially very good at, say, finding the recipe for a futuristic new material (or your email password). The classical machine is doomed to a life of stepwise calculation: Try one set of ingredients, fail, scrap everything, try again. But quantum computers can explore many potential recipes *simultaneously*.

So, naturally, tech giants such as Google, Huawei, IBM, and Microsoft have been chasing quantum's myriad positive applications—not only for materials science but also communications, drug development, and market analysis. China is plowing vast resources into state-backed efforts, and both the US and the European Union have pledged millions in funding to support homegrown quantum industries. Of course, whoever wins the race won't just have the next great engine of world-saving innovation. They'll also have the greatest code-breaking machine in history. So it's normal to wonder: What kind of Q-Day will humanity get—and is there anything we can do to prepare?

If you had a universal picklock, you might tell everyone—or you might keep it hidden in your pocket for as long as you possibly could. From a typical person's vantage point, maybe Q-Day wouldn't be recognizable as Q-Day at all. Maybe it would look like a series of strange and apparently unconnected news stories spread out over months or years. London's energy grid goes down on election day, plunging the city into darkness. A US submarine on a covert mission surfaces to find itself surrounded by enemy ships. Embarrassing material starts to show up online in greater and greater quantities: classified intelligence cables, presidential cover-ups, billionaires' dick pics. In this scenario, it might be decades before we're able to pin down exactly when Q-Day actually happened.

Then again, maybe the holder of the universal picklock prefers the disaster-movie outcome: everything, everywhere, all at once. Destroy the grid. Disable the missile silos. Take down the banking system. Open all the doors and let the secrets out.

Suppose you ask a classical computer to solve a simple math problem: Break the number 15 into its smallest prime factors. The computer would try all the options one by one and give you a near-instantaneous answer: 3 and 5. If you then ask the computer to factor a number with 1,000 digits, it would tackle the problem in exactly the same way—but the calculation would take millennia. This is the key to a lot of modern cryptography.

Take RSA encryption, developed in the late 1970s and still used for securing email, websites, and much more. In RSA, you (or your encrypted messaging app of choice) create a private key, which consists of two or more large prime numbers. Those numbers, multiplied together, form part of your public key. When someone wants to send you a message, they use your public key to encrypt it. You're the only person who knows the original prime numbers, so you're the only person who can decrypt it. Until, that is, someone else builds a quantum computer that can use its spooky powers of parallel computation to derive the private key from the public one—not in millennia but in minutes. Then the whole system collapses.

The algorithm to do this already exists. In 1994, decades before anyone had built a real quantum computer, an AT&T Bell Labs researcher named Peter Shor designed the killer Q-Day app. Shor's algorithm takes advantage of the fact that quantum computers run not on bits but on qubits. Rather than being locked in a state of 0 or 1, they can exist as both simultaneously—in superposition. When you run an operation on a handful of qubits in a given quantum state, you're actually running that same operation on those same qubits in *all* their potential quantum states. With qubits, you're not confined to trial and error. A quantum computer can explore all potential solutions simultaneously. You're calculating probability distributions, waves of quantum feedback that pile onto each other and peak at the correct answer. With Shor's algorithm, carefully designed to amplify certain mathematical patterns, that's exactly what happens: Large numbers go in one end, factors come out the other.

In theory, at least. Qubits are incredibly difficult to build in real life, because the slightest environmental interference can nudge them out of the delicate state of superposition, where they balance like a spinning coin. But Shor's algorithm ignited

STORIES & SOUNDS FROM AROUND THE WORLD

CONDÉ NAST TRAVELER

WOMEN WHO TRAVEL

LISTEN AND FOLLOW WHEREVER YOU GET YOUR PODCASTS

interest in the field, and by the 2010s, a number of projects were starting to make progress on building the first qubits. In 2016, perhaps sensing the nascent threat of Q-Day, the US National Institute for Standards and Technology (NIST) launched a competition to develop quantum-proof encryption algorithms. These largely work by presenting quantum computers with complex multidimensional mazes, called structured lattices, that even they can't navigate without directions.

In 2019, Google's quantum lab in Santa Barbara claimed that it had achieved "quantum supremacy." Its 53-qubit chip could complete in just 200 seconds a task that would have taken 100,000 conventional computers about 10,000 years. Google's latest quantum processor, Willow, has 105 qubits. But to break encryption with Shor's algorithm, a quantum computer will need thousands or even millions.

There are now hundreds of companies trying to build quantum computers using wildly different methods, all geared toward keeping qubits isolated from the environment and under control: superconducting circuits, trapped ions, molecular magnets, carbon nanospheres. While progress on hardware inches forward, computer scientists are refining quantum algorithms, trying to reduce the number of qubits required to run them. Each step brings Q-Day closer.

That's bad news not just for RSA but also for a dizzying array of other systems that will be vulnerable on Q-Day. Security consultant Roger A. Grimes lists some of them in his book *Cryptography Apocalypse*: the DSA encryption used by many US government agencies until recently, the elliptic-curve cryptography used to secure cryptocurrencies like Bitcoin and Ethereum, the VPNs that let political activists and porn aficionados browse the web in secrecy, the random number generators that power online casinos, the smartcards that let you tap through locked doors at work, the security on your home Wi-Fi network, the two-factor authentication you use to log in to your email account.

Experts from one national security agency told me they break the resulting threats down into two broad areas: confidentiality and authentication. In other words, keeping secrets and controlling access to critical systems. Chris Demchak, a former US Army officer who is a professor of cybersecurity at the US Naval War College and spoke with me in a personal capacity, says that a Q-Day computer could let an adversary eavesdrop on classified military data in real time. "It would be very bad if they knew exactly where all of our submarines were," Demchak says. "It would be very bad if they knew exactly what our satellites are looking at. And it would be very bad if they knew exactly how many missiles we had and their range." The balance of geopolitical power in, say, the Taiwan Strait could quickly tilt.

Beyond that real-time threat to confidentiality, there's also the prospect of "harvest now, decrypt later" attacks. Hackers aligned with the Chinese state have reportedly been hoovering up encrypted data for years in hopes of one day having a quantum computer that can crack it. "They wolf up everything," Demchak told me. (The US almost certainly does this too.) The question then becomes: How long will your sensitive data remain valuable? "There might be some needles in that haystack," says Brian Mullins, the CEO of Mind Foundry, which helps companies implement quantum technology. Your current credit card details might be irrelevant in 10 years, but your fingerprint won't be. A list of intelligence assets from the end of the Iraq War might seem useless until one of those assets becomes a prominent politician.

The threat to authentication may be even scarier. "Pretty much anything that says a person is who they say they are is underpinned by encryption," says Deborah Frincke, a computer scientist and national security expert at Sandia National Laboratories. "Some of the most sensitive and valuable infrastructure that we have would be open to somebody coming in and pretending to be the rightful owner and issuing some kind of command: to shut down a network, to influence the energy grid, to create financial disruption by shutting down the stock market."

The exact level of Q-Day chaos will depend on who has access to the first cryptographically relevant quantum computers. If it's the United States, there will be a "fierce debate" at the highest levels of government, Demchak believes, over whether to release it for scientific purposes or keep it secret and use it for intelligence. "If a private company gets there first, the US will buy it and the Chinese will try to hack it," she claims. If it's one of the US tech companies, the government could put it under the strict export controls that now apply to AI chips.

Most nation-state attacks are on private companies—say, someone trying to break into a defense contractor like Lockheed Martin and steal plans for a next-generation fighter jet. But over time, as quantum computers become more widely available, the focus of the attacks could broaden. The likes of Microsoft and Amazon are already offering researchers access to their primitive quantum devices on the cloud—and big tech companies haven't always been great at policing who uses their platforms. (The soldier who blew up a

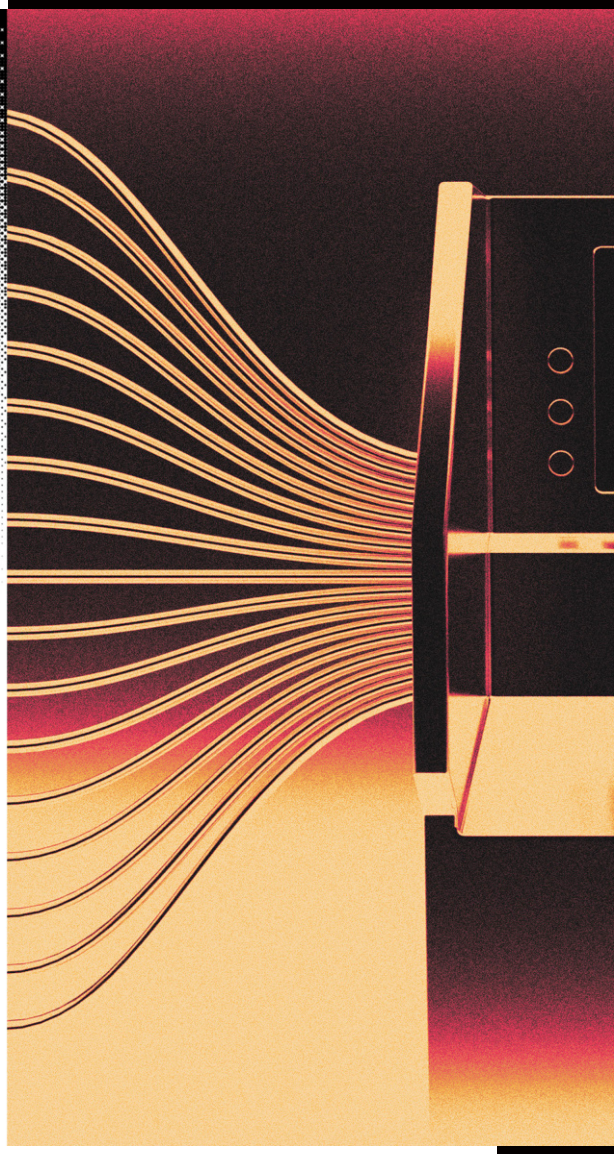

Cybertruck outside the Trump International Hotel in Las Vegas early this year queried ChatGPT to help plan the attack.) You could have a bizarre scenario where a cybercriminal uses Amazon's cloud quantum computing platform to break into Amazon Web Services.

Cybercriminals with access to a quantum computer could use it to go after the same targets more effectively, or take bigger swings: hijacking the SWIFT international payments system to redirect money transfers, or conducting corporate espionage to collect kompromat. The earliest quantum computers probably won't be able to run Shor's algorithm that quickly—they might only get one or two keys a day. But combining a quantum computer with an artificial intelligence that can map out an organization's weakness and highlight which keys to decrypt to cause the most damage could yield devastating results.

And then there's Bitcoin. The cryptocurrency is exquisitely vulnerable to Q-Day. Because each block in the Bitcoin blockchain captures the data from the previous block, Bitcoin cannot be upgraded to post-quantum cryptography, according to Kapil Dhiman, CEO of Quranium, a post-quantum blockchain security company. "The only solution to that seems to be a hard fork—give birth to a new chain and the old chain dies."

But that would require a massive organizational effort. First, 51 percent of Bitcoin node operators would have to agree. Then everyone who holds bitcoin would have to manually move their funds from the old chain to the new one (including the elusive Satoshi Nakamoto, the Bitcoin developer who controls wallets containing around $100 billion of the cryptocurrency). If Q-Day happens before the hard fork, there's nothing to stop bitcoin going to zero. "It's like a time bomb," says Dhiman.

That bomb going off will only be the beginning. When Q-Day becomes public knowledge, either via grim governmental address or cheery big-tech press release, the world will enter the post-quantum age. It will be an era defined by mistrust and panic—the end of digital security as we know it. "And then the scramble begins," says Demchak.

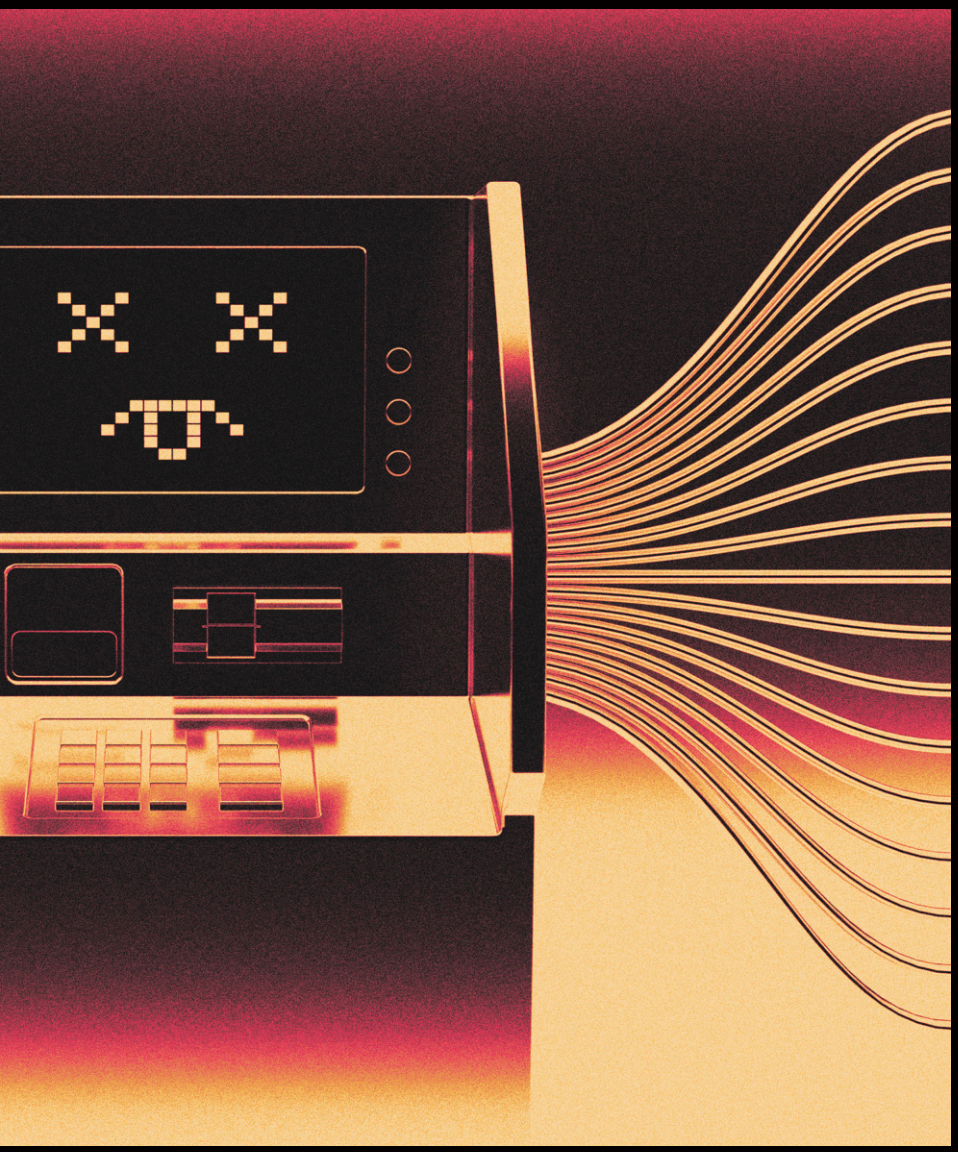

All confidence in the confidentiality of our communications will collapse. Of course, it's unlikely that everyone's messages will actually be targeted, but the perception that you could be spied on at any time will change the way we live. And if NIST's quantum-proof algorithms haven't rolled out to your devices by that point, you face a real problem—because any attempts to install updates over the cloud will also be suspect. What if that download from Apple isn't actually from Apple? Can you trust the instructions telling you to transfer your crypto to a new quantum-secure wallet?

Grimes, the author of *Cryptography Apocalypse*, predicts enormous disruptions. We might have to revert to Cold War methods of transmitting sensitive data. (It's rumored that after a major hack in 2011, one contractor purportedly asked its staff to stop using email for six weeks.) Fill a hard drive, lock it in a briefcase, put someone you trust on a plane with the payload handcuffed to their wrist. Or use one-time pads—books of pre-agreed codes to encrypt and decrypt messages. Quantum-secure, but not very scalable. Expect major industries—energy, finance, health care, manufacturing, transportation—to slow to a crawl as companies with sensitive data switch to paper-based methods of doing business and scramble to hire expensive cryptography consultants. There will be a spike in inflation. Most people might just accept the inevitable: a post-privacy society in which any

expectation of secrecy evaporates unless you're talking to someone in person in a secluded area with your phones switched off. Big Quantum is Watching You.

The best-case scenario looks something like Y2K, where we have a collective panic, make the necessary upgrades to encryption, and by the time Q-Day rolls around it's such an anticlimax that it becomes a joke. That outcome may still be possible. Last summer, NIST released its first set of post-quantum encryption standards. One of Joe Biden's last acts as president was to sign an executive order changing the deadline for government agencies to implement NIST's algorithms from 2035 to "as soon as practicable."

Already, NIST's post-quantum cryptography has been rolled out on messaging platforms such as Signal and iMessage. Sources told me that sensitive national security data is probably being locked up in ways that are quantum-secure. But while your email account can easily be Q-proofed over the internet (assuming the update doesn't come from a quantum imposter!), other things can't. Public bodies like the UK's National Health Service are still using hardware and software from the 1990s. "Microsoft is not going to upgrade some of its oldest operating systems to be post-quantum secure," says Ali El Kaafarani, the CEO of PQShield, a company that makes quantum-resistant hardware. Updates to physical infrastructure can take decades, and some of that infrastructure has vulnerable cryptography in places it can't be changed: The energy grid, military hardware, and satellites could all be at risk.

And there's a balance to be struck. Rushing the transition risks introducing vulnerabilities that weren't there before. "How do you make transitions slow enough that you can be confident and fast enough that you don't dawdle?" asks Chris Ballance, CEO of Oxford Ionics, a quantum computing company. Some of those vulnerabilities might even be there by design: Memos leaked by Edward Snowden indicate that the NSA may have inserted a backdoor into a pseudorandom number generator that was adopted by NIST in 2006. "Anytime anybody says you should use this particular algorithm and there's a nation-state behind it, you've got to wonder whether there's a vested interest," says Rob Young, director of Lancaster University's Quantum Technology Centre.

Then again, several people I spoke to pointed out that any nation-state with the financial muscle and technical knowledge to build a quantum device that can run Shor's algorithm could just as easily compromise the financial system, the energy grid, or an enemy's security apparatus through conventional methods. Why invent a new computing paradigm when you can just bribe a janitor?

Long before quantum technology is good enough to break encryption, it will be commercially and scientifically useful enough to tilt the global balance. As researchers solve the engineering challenge of isolating qubits from the environment, they'll develop exquisitely sensitive quantum sensors that will be able to unmask stealth ships and map hidden bunkers, or give us new insight into the human body. Similarly, pharma companies of the future *could* use quantum to steal a rival's inventions—or use it to dream up even better ones. So ultimately the best way to stave off Q-Day may be to share those benefits around: Take the better batteries, the miracle drugs, the far-sighted climate forecasting, and use them to build a quantum utopia of new materials and better lives for everyone. Or—let the scramble begin. W

AMIT KATWALA *is a features editor at* WIRED *and the author, most recently, of* Tremors in the Blood: Murder, Obsession, and the Birth of the Lie Detector.

THE NEW YORKER 100

Here's to the *next* hundred.

We just reached a milestone birthday. Explore a century of rigorous journalism, beloved cartoons, audio stories, games, and more. Subscribe today.

Scan to subscribe.

QUANTUM IS ALIVE!

Illustration by **Tal Uliel**

1981 In a lecture at MIT, the great physicist Richard Feynman proposes a powerful new computer built on quantum-mechanical principles. It could crack the hardest problems.

1985 Oxford astrophysicist David Deutsch describes the first universal quantum computer and proves it can simulate any physical system: It could, in short, put every classical computer to shame.

1992 Theoretical physicist Ben Schumacher coins the term "qubit" to describe the fundamental unit of a quantum computer: a transistor that can be 1, 0, or something in between.

1994 Deutsch and his colleague Richard Jozsa develop the first quantum algorithm that performs better than any possible classical algorithm on a given problem.

1996 At the annual Symposium on Foundations of Computer Science, Peter Shor debuts his quantum factoring algorithm. With just a few thousand qubits, quantum computers of the future will be able to break standard encryption! Quantum computers finally have a reason to exist.

1998 Lov Grover, a researcher at Bell Labs, develops another useful quantum algorithm. It can search unstructured databases—talk about practical!

2000 Researchers at IBM, MIT, and Berkeley build the first actual working quantum computer: a two-qubit device based on applying magnetic fields to a solution of chloroform molecules.

2011 Scientists show they can control a five-qubit quantum computer using nuclear magnetic resonance. The machine can run the Deutsch-Jozsa algorithm.

2016 A startup called D-Wave announces the world's first commercial quantum computer—a $10 million device designed to solve optimization problems.

2019 A new type of qubit emerges based on trapped ions. Researchers at the University of Innsbruck and MIT build a five-qubit quantum computer that can run Shor's algorithm.

2021 Google claims "quantum supremacy" with its 53-qubit Sycamore processor, performing a specific calculation faster than any classical computer. It uses a new approach based on superconducting qubits.

2024-2025 IBM breaks the 100-qubit barrier! Its previous processor, Hummingbird, had 65 qubits. This one, Eagle, has 127. (The following year's model, Osprey, will have 433.)

2025 Google launches Willow, a quantum chip with 105 qubits and improved error correction. It can do a task in five minutes that would take today's supercomputers 10 septillion years. Alphabet's stock price soars.

Oxford researchers announce a new method for creating high-fidelity connections between qubits, promising to cut down on decoherence and error.

Microsoft announces Majorana 1, a quantum chip powered by—it says—an entirely new state of matter. The company promises to bring full-scale quantum within reach in years rather than decades.

SCHRÖDINGER'S COMPUTER

This is the year of quantum science and technology. Literally: The United Nations said so. But no amount of pomp and hype can make something true. The fact is, it's been the year of quantum nearly every year for many, many years now. It's

QUANTUM IS DEAD.

1981 In a lecture at MIT, the out-there physicist Richard Feynman imagines a hypothetical quantum device. No one knows what problems it could solve.

1985 Oxford astrophysicist David Deutsch describes the first universal quantum computer, but good luck building it: It relies on pure quantum states that cannot be maintained in any lab environment.

1992 The term "qubit" is coined—but it's just a fancy bit of math. 1s and 0s at the same time? Even if that were possible, you'd need an entirely new class of algorithms.

1994 Behold the first quantum algorithm! It's an arbitrary math problem with zero practical application.

1996 Shor's famous algorithm—by which quantum computers can quickly factor large numbers—makes the first real-world case for quantum computers. But qubits are still struggling to exist in the real world: The slightest environmental interference can make them decohere back into 1 or 0.

1998 Grover's algorithm can search unstructured databases, but the speedup is only quadratic: Where a classical algorithm does something in 100 steps, Grover's can do it in 10. Not worth the engineering cost.

2000 The first working quantum computer has really high error rates. With more qubits added for error correction, you'd need millions of them to run Shor's algorithm.

2011 Scientists demonstrate a five-qubit quantum computer, but it's noisy: Any useful information gets drowned out. It'll never scale.

2016 A startup called D-Wave announces the world's first commercial quantum computer—but it uses something called quantum annealing, so it's not really a quantum computer. Researchers find no speed increase compared to classical machines.

2019 A five-qubit quantum computer based on trapped ions can run Shor's algorithm. But it can only factor the number 15. (You could do that on a pocket calculator in two seconds, or in your head in one.)

2021 Google's quantum supremacy breakthrough has no practical use, and IBM suggests that a classical computer could perform the same calculations in two and a half days (not the 10,000 years Google claimed).

2024-2025 IBM builds a chip with 100 qubits, but they decohere within microseconds. You'd need thousands more doing error correction. Plus, the giant fridges that cool the quantum chips down to near-absolute zero won't scale either.

2025 In a talk at CES, Nvidia's Jensen Huang says we're still decades away from useful quantum. Shares of quantum companies like IonQ, Rigetti, and D-Wave drop by many millions of dollars.

The AI boom threatens to steal quantum's thunder—not just by sucking up so much funding but also by stepping into its niche. There are fears of a quantum winter.

Quantum experts pour a bucket of cold water on Microsoft's claims of a new quantum breakthrough, pointing to a series of retractions from high-profile journals. The quantum uncertainty continues.

always the next big thing. It's always just about here. Well, is it? Really? Nearly a century ago, Erwin Schrödinger poked fun at quantum principles by comparing them to a cat that's both alive *and* dead. Today, quantum computing finds itself in just such a state: a superposition of here and not-here, a series of world-changing breakthroughs that are, at the same time, also colossal letdowns. So what's the real story? Let's open the box and find out. —THE EDITORS

ALL HAIL THE BASED GOD OF THERMODYNAMIC COMPUTING

Guillaume Verdon is building a new kind of chip to accelerate AI. His alter ego wants to accelerate humanity itself.

by **Will Knight**
Photographs by **Nicole Marie**

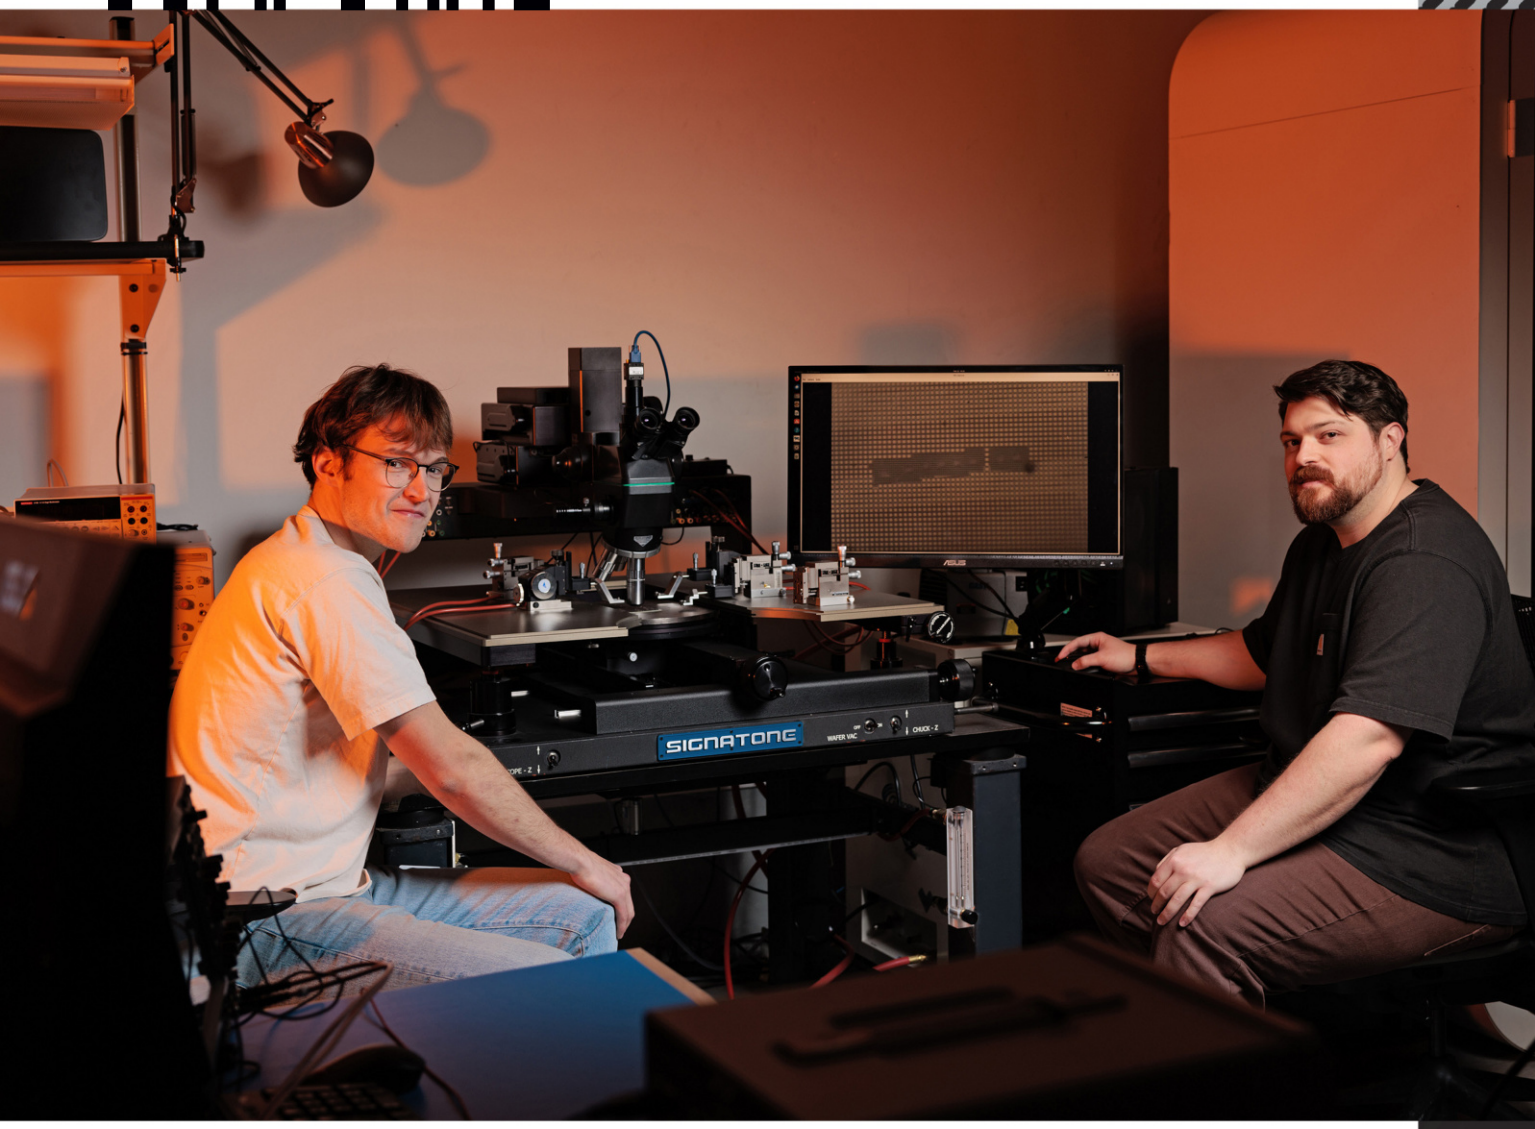

Guillaume Verdon stands before me with a new kind of computer chip in his hand—a piece of hardware he believes is so important to the future of humanity that he's asked me not to reveal our exact location, for fear that his headquarters could become the target of industrial espionage.

This much I can tell you: We're in an office a short drive from Boston, and the chip arrived from the foundry just a few days ago. It sits on a circuit board about the width of a Big Mac. The pinky-nail-sized piece of silicon itself is dotted with an exotic set of components: not the transistors of an ordinary semiconductor, nor the superconducting elements of a quantum chip, but the guts of a radically new paradigm called thermodynamic computing.

Not unlike its quantum cousin, thermodynamic computing promises to move beyond the binary constraints of 1s and 0s. But while quantum computing sets out—through extreme cryogenic cooling—to minimize the random

thermodynamic fluctuations that occur in electronic components, this new paradigm aims to *harness* those very fluctuations.

Engineers are chasing both paradigms in a race to accelerate past ordinary silicon chips and satisfy the ravenous demand for processing power in the age of AI. But Verdon—with his startup, Extropic—isn't just a contestant in that race. He's also one of the AI era's most shameless hype men. He is far better known as his online alter ego, Based Beff Jezos, the founding prophet of an ideology called effective accelerationism.

Known as "e/acc" for short, effective accelerationism is an irreverent rejection of effective altruism, a movement that has persuaded many technically minded people that the rise of artificial general intelligence—unless it is corralled and made safe—poses an almost certain existential risk to humanity. "EA's be like: 'I believe in Leprechauns and the burden of proof is on you to disprove me,'" went one fairly typical Based Beff post from 2022.

The AI existential risk movement, he wrote in another post that year, "is an infohazard that causes depression in our most talented and intelligent folks, killing our productive gains towards a greater more prosperous future." Another frequent target of his mockery is the AI ethics movement, which critiques large language models as riddled with the biases and blind spots of their architects. As Based Beff continued to spread e/acc's gospel online, it quickly became a rallying cry among some members of the tech elite, with prominent figures like Marc Andreessen and Garry Tan temporarily adding "e/acc" to their X usernames.

Effective accelerationism is perhaps best seen as the most technical fringe of a broader zeitgeist: a belief that American politics is broken and that caution, overregulation, and woke ideology are holding the country back. That ethos helped propel Donald Trump back to the White House, with Elon Musk, a hero of the e/acc movement, by his side. Like Trump 2.0, effective accelerationism promises an unstoppable American renaissance and an untroubled view of the work needed to get there. In both cases, the details are fuzzy.

But under the sharp resolution of a laboratory microscope, the specifics of Verdon's new chip are, if nothing else, plain to see: an array of square features each a few dozen microns wide. These components, Verdon promises, will be used to generate "programmable randomness"—a chip in which probabilities can be controlled to produce useful computations. When combined with a classical computer, he says, they will provide a highly efficient way to model uncertainty, a key task in all sorts of advanced computing, from modeling the weather and financial markets to artificial intelligence. (Some academic labs have already built prototype thermodynamic hardware, including a simple neural network—the technology at the heart of modern machine learning.)

When we meet at Extropic's headquarters in early 2025, Verdon, 32, looks more like a contractor than a physicist, wearing a plain Carhartt T-shirt over a broad frame, a pair of angular glasses, and a trim beard over a square jaw. Extropic, he explains, plans to have its first operational chips available for market later in 2025, less than a year after the design was first conceived. "Eventually entire workloads could run on thermo hardware," Verdon says. "Some day a whole language model could run on it."

The speed, maybe impatience, with which Verdon has developed these thermodynamic "accelerators" is all the more remarkable for another reason. Not long ago his work was synonymous with a completely different futuristic computing paradigm, only it wasn't moving fast enough for him. Effective altruism and AI safety aren't all Verdon has noisily repudiated over the past couple of years; he has also rejected quantum computing. "The story of thermodynamic computing is kind of an exodus from quantum," Verdon says. "And I kind of started that."

Guillaume Verdon (right) and his Extropic cofounder Trevor McCourt.

Why not save all that heat-killing energy and embrace the chaos—in hardware and in life—to make something new?

Back in 2019, when he was still a PhD student at Canada's University of Waterloo, Verdon was recruited to a team at Google tasked with figuring out how to use quantum computers—fabled machines capable of harnessing quantum mechanics to perform computations at unfathomable speeds—to supercharge artificial intelligence.

VOGUE
The World of Fashion at Your Fingertips
The New Vogue App Is Here
DISCOVER
PLAY
SHOP
DISCUSS
Download on the App Store
GET IT ON Google Play

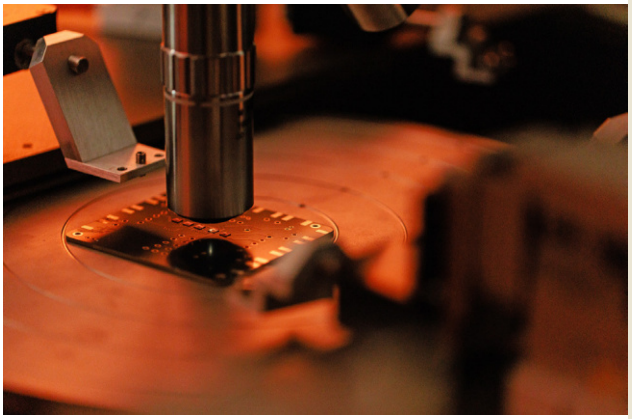

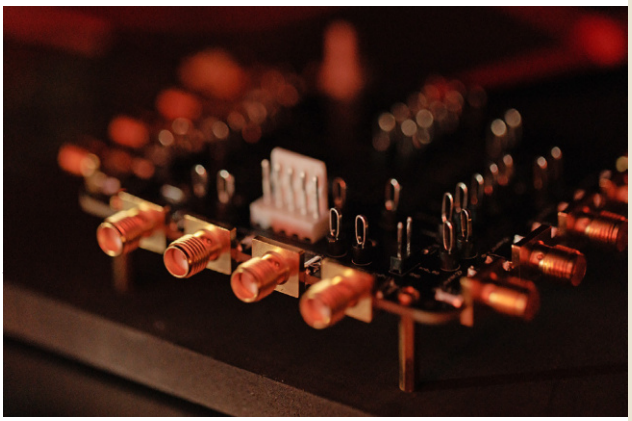

Top: Extropic's thermodynamic accelerator under a microscope. Bottom: The connectors that will integrate the chip into a classical computer.

By early 2022 the team had made important progress. Verdon had led the development of TensorFlow Quantum, a version of the company's AI software designed to run on quantum machines, and the hardware team at Alphabet was working on perfecting quantum error correction for a single qubit—a trick needed to account for the insane instability of quantum states.

But while Verdon was busy working at Google, several theoretical discoveries in academia started to suggest that quantum computing would take a lot longer to pay off than originally hoped. Back in 2018, Ewin Tang, then just an undergraduate student at the University of Texas at Austin, had published a startling paper showing that one of the best candidates for an exponential speedup in quantum machine learning—a kind of quantum recommendation algorithm—would not offer much acceleration after all. Other papers built upon the discovery and further impugned quantum's reputation.

Verdon was not the only person who felt that a quantum bubble was losing air. "It looked like all these applications we really wanted to do weren't going to work in the era of noisy quantum computing," says Faris Sbahi, cofounder and CEO of a thermodynamic startup called Normal Computing. "I became a little negative on the prospects for commercial quantum computing, and I switched fields." Both Normal Computing and Extropic have attracted other refugees from other quantum computing labs and startups. (Verdon says the two companies have different approaches and are not direct rivals.)

This quiet quantum bust coincided with the generative AI boom—and with another growing source of frustration for Verdon. Just as ChatGPT was taking off, Verdon felt that the industry's response was dampened by effective altruists who wanted to freeze progress and subject the technology to needless controls. He saw the tech industry itself consumed by an overly scrupulous, self-critical, woke ideology.

Quantum dead ends, tech-industry navel-gazing, an AI gold rush: All of these merged together in Verdon's mind to inspire a new philosophy along with belief in a new engineering idea. In both conventional and quantum computers, heat is a source of errors. Both kinds of computing require huge expenditures of energy to prevent natural thermodynamic effects in electronic circuits from either flipping 1s to 0s or destroying delicate quantum states. Instead of fighting these effects, Verdon thought, why not save all that heat-killing energy and embrace the chaos—in hardware and in life—to make something new?

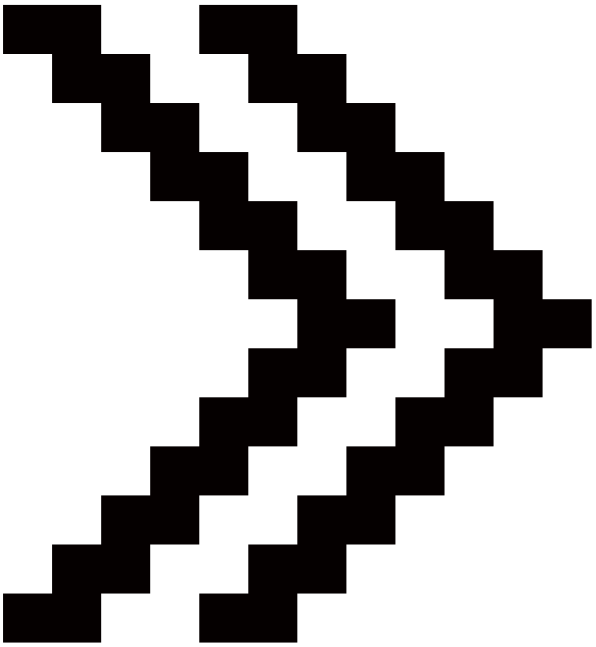

In *Lord of Light*, a 1967 science fiction novel by Roger Zelazny, the last member of a futuristic revolutionary group known as the accelerationists wants to convince society to embrace unrestrained technological progress. By the 1990s, a British philosopher named Nick Land was advocating for a real accelerationist movement that would unshackle capitalism from the restraints imposed by politicians and welcome the technological and social destruction and renewal this would bring. Accelerationist ideas are echoed by other alt-right thinkers, including the influential blogger Curtis Yarvin, who argues that Western democracy is a bust and ought to be replaced.

Verdon and his bros, however, see accelerationism through the prism of natural laws. In a founding manifesto, offering a "physics-first view of the principles underlying effective accelerationism," they cite recent work suggesting that the existence of life itself may be explained by the propensity of matter to harness and exploit energy. Not only does the thermodynamic approach promise a much leaner form of computing, Verdon argues, it also reflects how biological intelligence itself has evolved. He and his e/acc allies argue that the same energy-harnessing principle accounts for the development of social organizations and the very course of humanity.

"Effective accelerationism aims to follow the 'will of the universe': leaning into the thermodynamic bias towards futures with greater and smarter civilizations that are more effective at finding/extracting free energy from the universe and converting it to utility at grander and grander scales," the manifesto reads. And it doesn't hold humanity particularly dear: "e/acc has no particular allegiance to the biological substrate for intelligence and life."

In January 2022, while he was still at Google, Verdon decided to create a new X account, to talk shit and air some of his ideas. But he didn't feel free to do it under his own name. ("My tweets were closely monitored," he tells me.) So he called himself Based Beff, with a profile picture of a ripped '80s video game version of Jeff Bezos.

Based Beff quickly demonstrated a talent for humor, whipping up like-minded, right-leaning tech bros and trolling AI doomers. In July, he linked to an erotic novel about Microsoft's Clippy—and used it to poke fun at a canonical thought experiment, famous among effective altruists, involving a world-destroying AI that is optimized for making paper clips. "It's #PrimeDay so I'm buying this for all the anti-AGI EAs," he wrote. "I found it very freeing," Verdon says. "It was like a video game persona, and stuff just started pouring out of me."

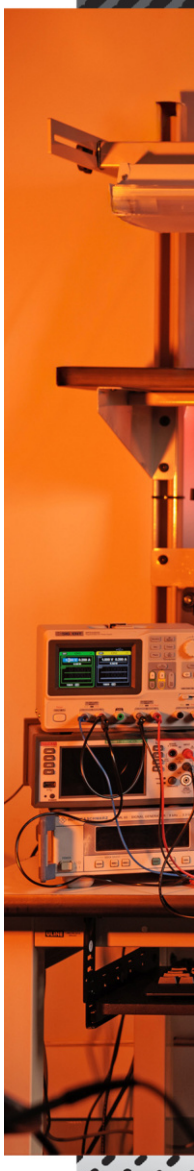

I first spoke to Verdon in December 2023, just after Forbes had used his voiceprint to identify him as the person behind the Based Beff Jezos account. The attention was uncomfortable for someone accustomed to working in relative obscurity. "I was in shock for a while," Verdon says. "It was traumatic."

But rather than run and hide after being outed, Verdon chose to embrace his newfound celebrity and use it to promote his company. Extropic hurriedly came out of stealth, revealed its vision, and set about recruiting. The company has so far raised $14.1 million in seed funding, and there are around 20 engineers working full-time at the firm. Investors include Garry Tan, Balaji Srinivasan, and two authors of the famous paper "Attention Is All You Need," which laid out the so-called transformer architecture that has revolutionized AI.

Verdon is a lot more chill in person than he is when posting as Based Beff Jezos. After I arrive at Extropic, Verdon makes me a very strong espresso before leading me to a windowless hardware lab where a handful of engineers are staring at diagrams of the company's first chip, fiddling with oscilloscopes that reveal performance characteristics, and staring at code designed to make it sing.

Verdon introduces me to his CTO and cofounder, Trevor McCourt, a tall, easygoing fellow with shaggy hair who also worked on quantum computing at Alphabet before becoming disillusioned with the project. I ask if McCourt, too, has an online alter ego. "No," he laughs. "One of me is enough."

A couple of weeks earlier, Google had revealed its

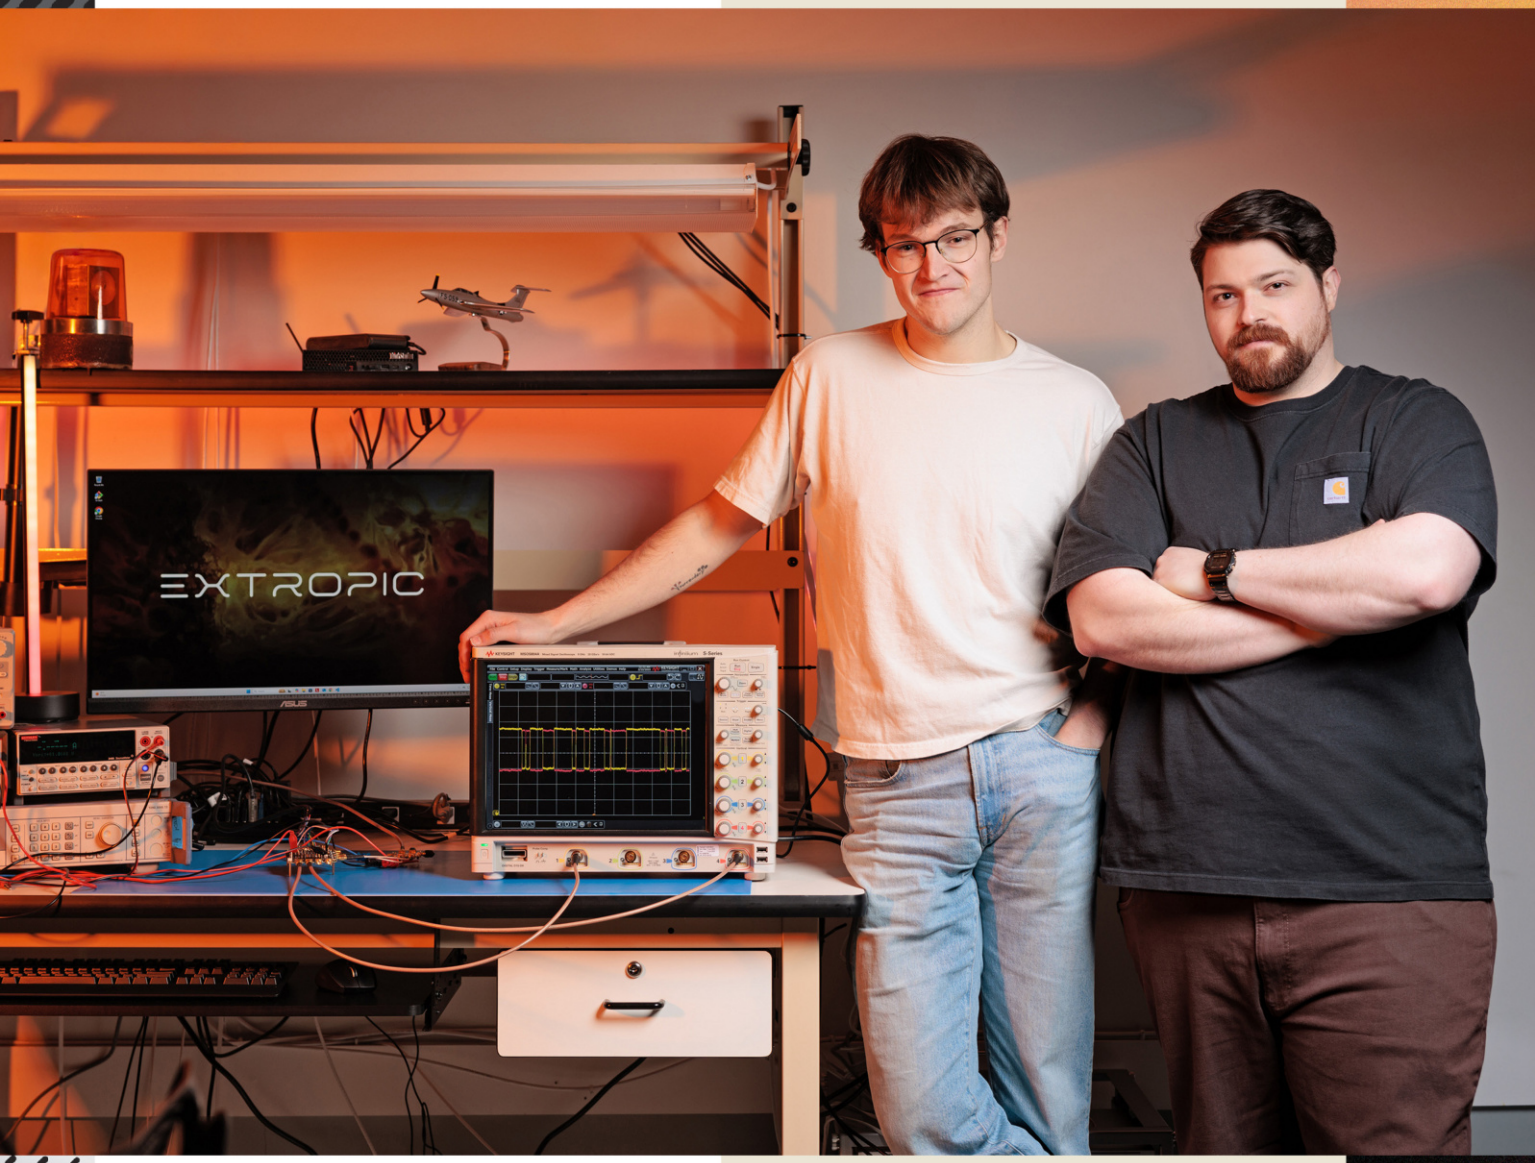

An oscilloscope displays the thermodynamic output—what Extropic calls a "probabilistic bit"—from McCourt and Verdon's chip.

He called himself Based Beff, with a profile picture of a ripped '80s video game version of Jeff Bezos.

latest quantum computing chip, called Willow. The chip can now perform quantum error correction on the scale of 105 qubits. For Extropic, this is not a sign that quantum is now on the fast track. Verdon and McCourt point out that the machine still needs to be cooled to extremely low temperatures, and its advantages remain uncertain.

Once early adopters get their hands on Extropic's hardware later this year, the chip may prove its worth with high-tech trading and medical research—both use algorithms that run probabilistic simulations. But in the meantime, the day after we meet, Verdon is on his way to DC for an event timed to Donald Trump's inauguration. Before boarding his flight, he makes a flurry of posts on X. "I love the techno-capital machine and its creations," he writes. "We're accelerating to the stars."

WILL KNIGHT *is a senior writer at* WIRED *covering artificial intelligence.*

64,000 STRAWBERRIES

THE WEIGHT

The internet is massive. But does it have … actual mass? Big server farms and miles of fiber-optic cables do, of course, but we don't mean the infrastructure of the internet. We mean the internet itself. The information. The data. The cybernetics. And because storing and moving stuff through cyberspace requires energy—which, per Einstein, has mass—it should, in theory, be possible to calculate the internet's weight.

Way back in the adolescent days of the web, in 2006, a Harvard physicist named Russell Seitz made an attempt. His conclusion? If you consider the mass of the energy powering the servers, the internet comes out to roughly 50 grams—or about the weight of a couple strawberries. People still use Seitz's comparison to this day. *We're all wasting our lives on something we could swallow in one bite!*

But a lot has happened since 2006—Instagram, iPhones, and the AI boom, to name a few. (By Seitz's logic, the internet would now weigh as much as a potato.) There's also the fact that, around the time of Seitz's calculation, Discover magazine proposed a different method. Information on the internet is written in bits, so what if you looked at the weight of the electrons needed to encode those bits? Using all internet traffic—then estimated to be 40 petabytes—Discover put the internet's weight at a tiny fraction (5 millionths) of a gram. So, more like a squeeze of strawberry juice. WIRED thought it was time to investigate for ourselves.

First up: the server-energy method. "Fifty grams is just wrong," says Christopher White, president of NEC Laboratories America and a veteran of storied research powerhouse Bell Labs. Other scientists we spoke to agreed. Daniel Whiteson, a particle physicist at UC Irvine and cohost of the podcast *Daniel and Kelly's Extraordinary Universe*, said it's an overly convenient way to get "the units you want"—like assuming the price of a doughnut could be calculated by dividing the total number of doughnuts in the world by the world GDP. Preposterous! That would give us a doughnut-per-dollar figure, sure, "but it wouldn't be correct, or even close," Whiteson says.

GETTY IMAGES

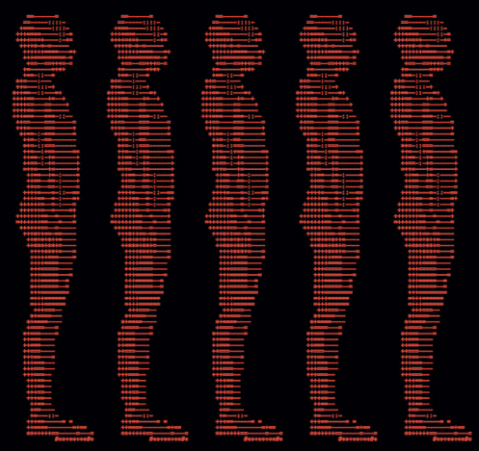

RELEVANT EQUATIONS + NUMBERS

$E = mc^2$

$P = VA$

LANDAUER'S PRINCIPLE

INTERNET IN 2025 = 175 ZETTABYTES

1G DNA = 215 PETABYTES

SAMANTHA SPENGLER *is a research editor at* WIRED.

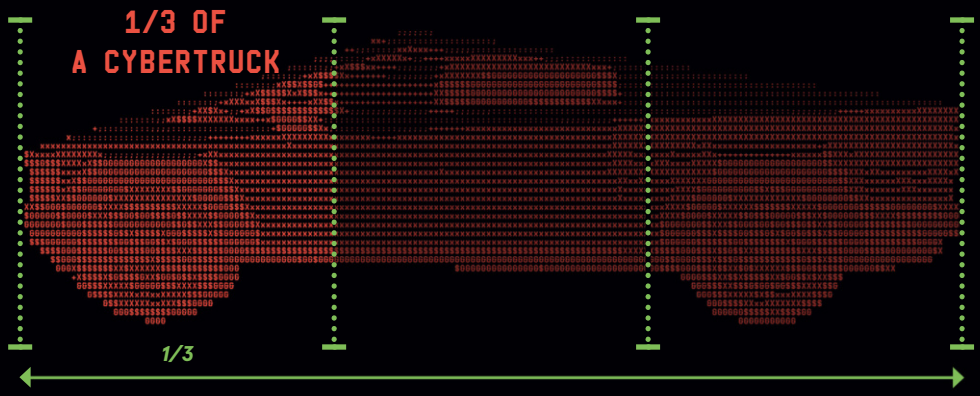

OF THE INTERNET

A MICROINVESTIGATION BY SAMANTHA SPENGLER

Discover magazine's calculation also seemed a little off to us. It has more to do with the transmission of the internet, as opposed to the internet itself. It also assumes a set number of electrons needed to encode information. In reality, the number is incredibly varied and depends on the specific chips and circuits being used.

White suggested a third method. What if we pretend to put all the data stored on the internet, across all the hundreds of millions of servers around the world, in just one place? How much energy would we need to encode that data, and how much would that energy weigh? In 2018, the International Data Corporation estimated that by 2025, the internet's datasphere would reach 175 zettabytes, or 1.65×10^{24} bits. (1 zettabyte = 1024^7 bytes and 1 byte = 8 bits.) White suggested multiplying those bits by a mathematical term—$k_B T \ln 2$, if you're curious—that captures the minimum energy needed to reset a bit. (Temperature is a factor, because storing data is easier in colder conditions. Meaning: The internet is lighter in space than it is in Tucson, Arizona.) We can then take that number, which will represent energy, and call on $E = mc^2$ to reach the total mass. At room temperature, the entirety of the internet would weigh $(1.65 \times 10^{24}) \times (2.9 \times 10^{-21})/c^2$, or 5.32×10^{-14} grams. That's 53 *quadrillionths* of a gram.

Which ... is no fun. Even if it has almost no physical mass, the internet still *feels* weighty, to those billions of us weighed down by it every day. White, who has previously attempted similar philosophical estimates, clarified that in reality, the web is so intricate that it is "essentially unknowable," but why not try? In recent years, scientists have floated the idea of storing data within the building blocks of nature: DNA. So what if we were to weigh the internet in those terms? Current estimates say that 1 gram of DNA can encode 215 petabytes—or 215×10^{15} bytes—of information. If the internet is 175×1024^7 bytes, that's 960,947 grams' worth of DNA. That's the same as 10.6 American males. Or one third of a Cybertruck. Or 64,000 strawberries. W

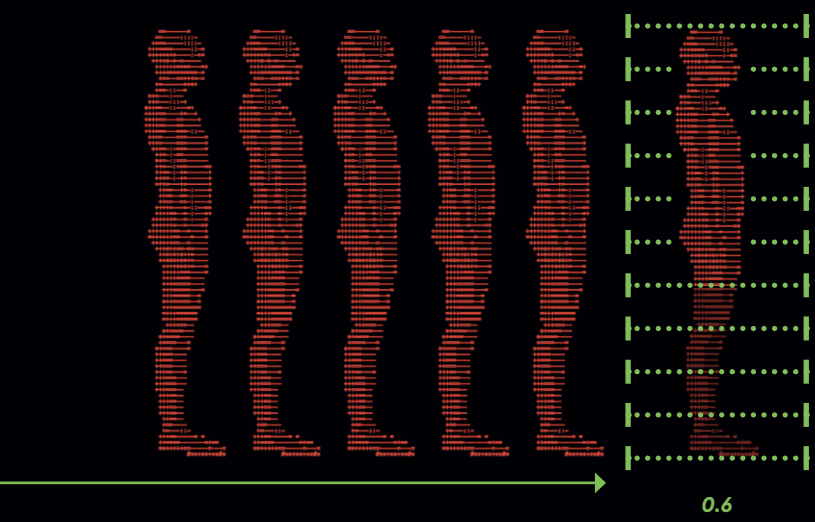

ANGELINA JOLIE WAS

Thirty years ago, she made an absurdly geeky prediction about the inner workings of digital architecture. It came true—twice.

RIGHT ABOUT COMPUTERS

by Jason Kehe

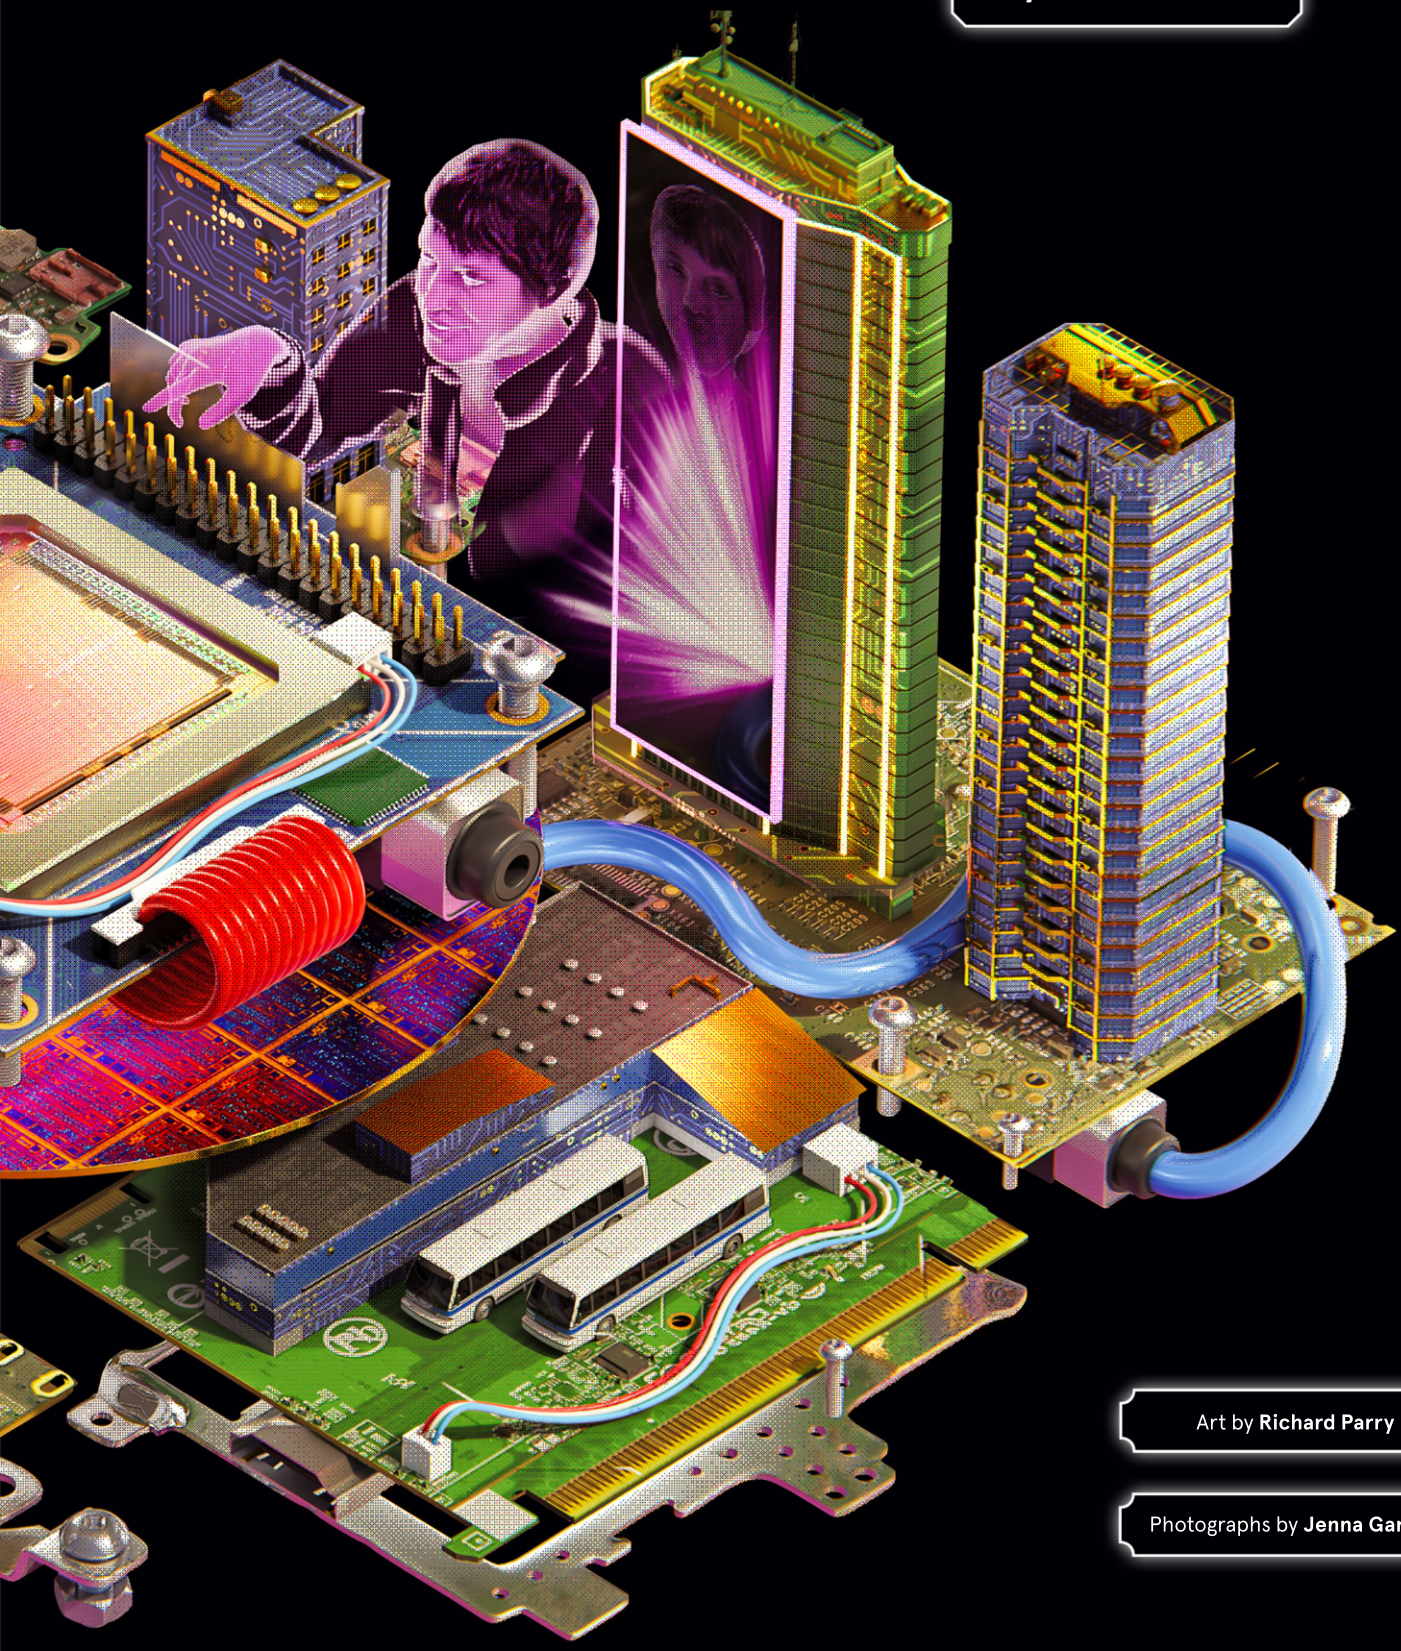

Art by **Richard Parry**

Photographs by **Jenna Garrett**

Incredibly, Angelina Jolie called it. The year was 1995. Picture Jolie, short of both hair and acting experience, as a teenage hacker in *Hackers*. Not a lot of people saw this movie. Even fewer appreciated its relevance. *Hackers* was "grating," Entertainment Weekly huffed at the time, for the way it embraced "the computer-kid-as-elite-rebel mystique currently being peddled by magazines like WIRED." Thirty years later, Entertainment Weekly no longer publishes a magazine, WIRED does, and *Hackers* ranks among the foundational documents of the digital age. The last time I saw the movie, it was being projected onto the wall of a cool-kids bar down the street from my house.

But that's not the incredible thing. The incredible thing, again, is that Jolie called it. *It.* The future. Midway through *Hackers*, she's watching her crush (played by Jonny Lee Miller, whom she'd later marry in real life) type passionately on a next-gen laptop. "Has a killer refresh rate," Miller says, breathing fast. Jolie replies: "P6 chip. Triple the speed of the Pentium." Miller's really worked up now. Then Jolie leans forward and, in that come-closer register soon to make her world-famous, says this: "RISC architecture is gonna change everything."

You have to believe me when I say, one more time, that this is incredible. And what's incredible is not just that the filmmakers knew what RISC architecture was. Or that Jolie pronounced it correctly ("risk"). Or even that Jolie's character was right. What's incredible is that she's *still* right—arguably even *more* right—today. Because RISC architecture is, somehow, changing everything again, here in the 21st century. Who makes what. Who controls the future. The very soul of technology. *Everything.*

And nobody's talking about it.

And that's probably because the vast majority of people everywhere, who use tech built on it every single day, still don't know what in the computer-geek hell a RISC architecture even is.

Unless you're *in* computer-geek hell, as I am, right now. I've just arrived at the annual international RISC-V (that's "risk five") summit in Santa Clara, California. Here, people don't just know what RISC is. They also know what, oh, vector extensions and AI accelerators and matrix engines are. At the coffee bar, I overhear one guy say to another: "This is a very technical conference. This is a very technical *community*." To which the other guy replies: "It ought to be. It ought to be."

OK, but where are the cool kids? It's hard not to fixate on appearances at an event like this—a generic convention center, with generic coffee, in a generic town. I guess I was hoping for neon lights and pixie cuts. Instead it's frumpy, forgettable menswear as far as the eye can see. There are 30 men for every woman, I count, as everyone gathers in the main hall for the morning presentations.

Then someone takes the stage, and she's not just a she. She is Calista Redmond, the CEO of RISC-V International, and, Angelina Jolie be praised, she's wearing a nifty jacket, a statement belt, and gold-and-silver … pumps? stilettos? Wait, what's the difference? Of all the things to ask Redmond when I run into her at a happy hour later that day, that's what I choose. She looks at me, smiles blankly, and just says, "I don't know."

In shame I retreat to the bar, where I decide I must redeem myself. So, cautiously, I make my way back to Redmond, who's now deep in conversation with the chief marketing officer of a semiconductor startup. I try to impress them with a technical observation, something about RISC and AI. Redmond turns to me and says, "I thought you wanted to talk about shoes." I assure her I'm not here to talk about what's on the outside. I'm here to talk about what's on the inside.

"Jason here is writing a story about RISC for WIRED," Redmond tells the CMO. She's not sure, frankly, that this is a great idea. Not because she isn't a believer. In many ways, she's *the* believer, the face of the brand. Attendees at the conference invoke her name with casual reverence: *Calista*

says this, Calista thinks that. And did you hear her morning keynote? In fact I did. "We have fundamentally launched!" she announced, to the yelps of the business-casuals. RISC-V will transform, is transforming, machinery everywhere, she said, from cars to laptops to spaceships. If anyone doubts this, Redmond sends them the *Hackers* clip.

So why, I press her now, should I not support the cause and write the big, cyberpunky, untold story of RISC? Because, Redmond says, not only does no one know what RISC is. No one *cares* what RISC is. And no one should. People don't buy "this or that widget," she says, because of what's inside it. All they want to know is: Does the thing work, and can I afford it?

To my dismay, almost everyone I talk to at the conference agrees with Redmond. Executives, engineers, marketers, the people refilling the coffee: "Calista's probably right," they say. Now it's my turn to get annoyed. I thought insides mattered! RISC is one of the great and ongoing stories of our time! People should care.

So I resolve to talk to the one person I think *must* agree with me, who has to be on my side: the legendary inventor of RISC itself.

The inner workings of a computer, David Patterson says, should be kept simple, stupid. We're sitting in an engineering lab at UC Berkeley, and Patterson—77 years old, partial to no-frills athleisure—is scribbling on a whiteboard. A computer's base operation, he explains, is the simplest of all: ADD. From there you can derive SUBTRACT. With LOAD and STORE, plus 30 or so other core functions, you have a complete basis for digital computation. Computer architects call this the "instruction set architecture," or the ISA. (They switch between saying each letter, "I-S-A," and—the neater option—pronouncing it as a word, "eye-suh.")

Computer architectures are so named because, well, that's exactly what they are—architectures not of bricks but of bits. The people who made *Hackers* plainly understood this. In sequences of dorky-awesome special effects, we fly through futuristic streets, look up at futuristic buildings, only to realize: This isn't a city. This is a microchip.

Even within a chip, there are subarchitectures. First come the silicon atoms themselves, and on top of those go the transistors, the circuits and gates, the microprocessors, and so on. You'll find the ISA at the highest layer of the hardware. It is, I think, the most profound architecture ever devised by humans, at any scale. It runs the CPU, the computer's brain. It's the precise point, in other words, at which dead, inert, hard silicon becomes, via a set of powerful animating conjurations, soft and malleable—alive.

Everyone has their own way of explaining it. The ISA is the *bridge*, or the *interface*, between the hardware and the software. Or it's the *blueprint*. Or it's the computer's *DNA*. These are helpful enough, as is the common comparison of an ISA to a *language*. "You and I are using English," as Redmond said to me at the conference. "That's our ISA." But it gets confusing. Software speaks in languages too—programming languages. That's why Patterson prefers *dictionary* or *vocabulary*. The ISA is less a specific language, more a set of generally available words.

Back when Patterson started out, in the 1970s, the early ISAs were spinning out of control. Established tech companies figured that as hardware design improved and programming languages got more sophisticated, computers shouldn't remain

EVERETT COLLECTION

In *Hackers*, a group of RISC-obsessed teens must stop a goateed villain from capsizing a fleet of oil tankers with a computer virus. Pictured here, from left: Lord Nikon, Crash Override (aka Zero Cool), Acid Burn, and Cereal Killer.

simple; they should be taught *larger* vocabularies, with *longer* words. The more types of operations they were capable of, the logic went, the more efficient their calculations would be.

On the whiteboard, Patterson scrawls the word POLYNOMIAL in big letters—just one of the hundreds of operations that Intel and others added to their ISAs. Even as a young recruit at Berkeley, Patterson suspected that the bigwigs had it backward, that exactly none of these esoteric add-ons were necessary. That a bigger dictionary did not lead to clearer sentences.

So he and a senior colleague decided to strip the kruft from the instruction sets of midcentury computing. At the time, the Defense Advanced Research Projects Agency was giving out grants for "high-risk" research. Patterson says they chose the acronym RISC—*reduced* instruction set computer—as a fundraising ploy. Darpa gave them the money.

Patterson then did as aspiring academics do: He wrote a spicy paper. Called "The Case for the Reduced Instruction Set Computer" and published in 1980, it set off a great war of architectures. "The question then," as Patterson would later say in an acceptance speech for a major prize, "was whether RISC or CISC was faster." CISC (pronounced "sisk") was the name Patterson gave the rival camp: *complex* instruction set computer. The CISCites fired back with a paper of their own and, at international conferences throughout the early '80s, battled it out with the RISCites onstage, the bloodshed often spilling into the hallways and late-night afterparties. Patterson taunted his opponents: They were driving lumbering trucks while he was in a feather-light roadster. If you magnify a RISC-based microchip from those years, you'll spot a sports car etched into the upper left corner, just 0.4 millimeters in length.

The RISCites won. With vigilant testing, they proved that their machines were between three and four times faster than the CISC equivalents. The RISC chips had to perform more operations per second, it's true—but would you rather read a paragraph of simple words, or a sentence of polysyllabic verbiage? In the end, CISCites retracted their claims to supremacy, and the likes of Intel turned to RISC for their architectural needs.

Not that anybody outside tech circles talked about this at the time. When *Hackers* came out in 1995, Patterson was flabbergasted to hear his life's work, 15 years old by that point, mentioned so casually and seductively by a Hollywood starlet. Computers were still too geeky, surely, to matter to the masses. (When I make Patterson rewatch the scene, he's all smiles and pride, though he does say they mistake "refresh rate" for "clock rate.")

Still, Patterson's invention was indeed chang-

Computer architecture isn't telling a machine what to do. It's establishing the possibility that it can be told anything at all. The work is superhuman, if not fully alien.

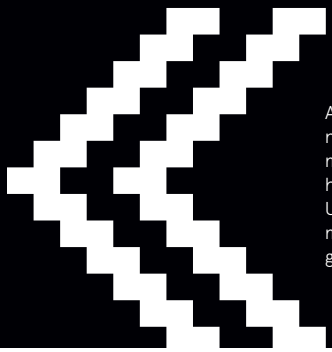

As CEO of RISC-V International, Calista Redmond moved the foundation's headquarters from the US to Switzerland to allay members' concerns about geopolitical neutrality.

ing everything. In those years, a rising company in the UK called Arm—the "r" in its name stood for RISC—was working with Steve Jobs on tablet-sized devices that needed smaller, faster CPUs. That effort stalled, but one thing led to another, and if you're reading this on a phone right now, you have RISC-based Arm architectures to thank. When Patterson walks me out of the Berkeley building at the end of our dizzying afternoon together, we stop by a handsome bronze plaque in the lobby that commemorates his "milestone" creation of the first RISC microprocessor. We stare at it in prayerful awe. "1980–1982," it reads—the bloodiest years of the great architecture war.

Better make room for another plaque, I note.

The year is now 2008. Two instruction sets exert near-total control over digital life. One is called x86, the descendent of Intel's legacy CISC architecture, and it dominates the high end of machinery: personal computers and servers. Arm's RISC architecture, meanwhile, dominates everything else: phones, game consoles, the internet of things. Different though they are, and with opposite origins, these two ISAs share one important feature: They're both *closed*, proprietary. You can't modify them, and if you want to use them, you have to pay for them.

Andrew Waterman, a graduate student at—where else?—UC Berkeley, finds this frustrating. As a computer architect, he wants to build things, deep things. Things at the very foundations of computing. But right now he has no good ISAs to play with. Arm and x86 are off-limits, and the free architectures for students are just so ... baggy. They use register windows to speed up procedure calls, for God's sake! Never mind what that means. The point is, every person in this story is a genius.

So Waterman and two other geniuses have an idea: Why not create a new, better-working, free ISA for academic use? It's an idea they know someone else has had before. To Patterson they go. And because he's their inspiration, and because he has worked on four generations of RISC architectures by this point, they'll call it, they announce to him proudly, RISC-V. Patterson is touched. A bit skeptical, sure, especially when they say they'll be done in three months. But touched. He gives the boys his blessing, his resources, and a classic bit of advice: Keep it simple, stupid.

RISC-V does not take three months. It takes closer to four years. If I've failed, so far, to account for the precision of this work, let me try again here. Computer architects are not software engineers, who use programming languages to talk to the machine. Even coders who can speak assembly or C, the so-called low-level languages, still do just that: They talk. Computer architects need to go deeper. Much deeper. All the way down to a preverbal realm. If they're speaking at all, they're speaking in gestures, motions: the way primitive circuits hold information. Computer architecture isn't telling a machine what to do. It's establishing the possibility that it can be told anything at all. The work is superhuman, if not fully alien. Put it this way: If you found the exact place in a human being where matter becomes mind, where body becomes soul—a place that no scientist or philosopher or spiritual figure has found in 5,000 years of frantic searching—wouldn't you tread carefully? One wrong move and everything goes silent.

In 2011, Waterman and his two collaborators, Krste Asanović and Yunsup Lee, release RISC-V into the wild. They've accomplished their mission: Geeky grad students everywhere, and hobbyists too, have an ISA for whatever computer-architecting adventures they might undertake. These early days feel utopian. Then Patterson, a proud dad, does as retiring academics do: He writes a spicy paper. Called "The Case for Open Instruction Sets" and published in 2014, it sets off a—

Yes. We've been here before. A second war of the architectures.

It's hard to overstate just how topsy-freaking-turvy this gets. To review: Patterson invented RISC in 1980 and went to battle with the established ISAs. He won. Thirty years later, his disciples reinvent RISC for a new age, and he and they go to battle with the very company whose success secured RISC's legacy in the first place: Arm.

In response to Patterson's paper, Arm fires back with a rebuttal, "The Case for Licensed Instruction Sets." Nobody wants some random, untested, unsupported ISA, they say. Customers want success, standards, a proven "ecosystem." The resources it would take to retool and reprogram everything for a new ISA? There's not enough cash in the world, Arm scoffs.

The RISC-V community disagrees. They create their own ecosystem under the auspices of RISC-V International and begin adapting RISC-V to the needs of modern computing. Some supporters start calling it an "open source hardware" movement, even if hardcore RISC-Vers don't love the phrase. Hardware, being set in literal stone, can't

THE ORIGINAL 37 INSTRUCTIONS OF THE RISC-I ARCHITECTURE (ORDERED BY TYPE):

ADD (REGISTER/IMMEDIATE)
ADD WITH CARRY
SUBTRACT (REGISTER/IMMEDIATE)
SUBTRACT WITH CARRY
SUBTRACT REVERSE
SUBTRACT REVERSE WITH CARRY
EXCLUSIVE OR (REGISTER/IMMEDIATE)
OR (REGISTER/IMMEDIATE)
AND (REGISTER/IMMEDIATE)
LOAD BYTE SIGNED
LOAD SHORT SIGNED
LOAD LONG
LOAD BYTE UNSIGNED
LOAD SHORT UNSIGNED
STORE BYTE
STORE SHORT
STORE LONG
SHIFT LEFT LOGICAL (REGISTER/IMMEDIATE)
SHIFT RIGHT LOGICAL (REGISTER/IMMEDIATE)
SHIFT RIGHT ARITHMETIC (REGISTER/IMMEDIATE)
LOAD IMMEDIATE HIGH
SUBTRACT, SET CONDITION CODES, R0 DESTINATION
JUMP PC-RELATIVE ON EQUAL
JUMP PC-RELATIVE ON NOT EQUAL
JUMP PC-RELATIVE ON LESS THAN
JUMP PC-RELATIVE ON GREATER THAN OR EQUAL
JUMP PC-RELATIVE ON LESS THAN UNSIGNED
JUMP PC-RELATIVE ON GREATER THAN OR EQUAL UNSIGNED
CALL PC-RELATIVE
CALL REGISTER

exactly be "open source," and besides, RISC-V doesn't count, entirely, as hardware. It's the hardware-software *interface*, remember. But, semantics. The point stands: Anyone, in any bedroom or garage or office in any part of the world, can use RISC-V for free to build their own computers from scratch, to chart their own technological destiny.

Arm is right about one thing, though: This does take money. Millions if not billions of dollars. (If you think "fabless" chip printers can do it for closer to five figures, come back to me in five years.) Still, RISC-V begins to win. Much as Arm, in the 1990s and 2000s, found success in low-end markets, so too, in the 2010s, does RISC-V: special-purpose gadgets, computer chips in automobiles, that sort of thing. Why pay for Intel chips or Arm licenses when you don't have to?

And the guys at Berkeley? In 2015, they launch their own company, called SiFive, to build computer parts based on RISC-V. Meaning: Arm isn't just a spiritual enemy for them now. It's a direct competitor.

By the time I went to that "very technical conference" in Santa Clara, the Arm-vs.-RISC-V war had been raging for nearly a decade. I could still feel it everywhere. *We've won,* I heard several times. *Nobody's happy at Arm,* someone claimed. (One longtime higher-up at Arm, who insisted on anonymity to discuss internal affairs, disputed "nobody" but admitted there's been a "culture change" in recent years.) On the second day of the conference, when news broke of a rift between Arm and one of its biggest customers, Qualcomm, people cheered in the hallways. "Arm is assholes," a former SiFive exec told me. In fact, only one person at the conference seemed to have anything nice to say about the competition. He was working a demo booth, and when I marveled that his product was built on a RISC-V processor, he turned a little green and whispered: "Actually, it's Arm. Don't tell anyone. Please don't tell anyone."

Booth bro was probably worrying too much. In the hardware world, everyone has worked, or has friends, everywhere else. Calista Redmond, the star of the show, spent 12 years at IBM (and recently resigned from RISC-V International for a job at Nvidia). Even Patterson has ties to, of all places, Intel—which, though less of a direct threat to Arm, is still a RISC-V competitor. It was Intel grant money, Patterson happily admits, that paid for the Berkeley architects to invent RISC-V in the first place. Without closed source, proprietary Big Tech, there's no open source, free-for-all Little Tech. Don't listen to the techno-hippies who claim otherwise; that's always been the case.

Patterson was the big-ticket speaker on the second day of the conference, and in his talk, he brought up the paper that Arm wrote in rebuttal to his, lo those 10 years ago. One of its two authors has since parted ways with Arm. The other, Patterson noted, not only left—he now works at SiFive. "It's satisfying," Patterson said, "he has come to his senses." Which got a laugh, of course, but I was still stuck on something Patterson said earlier in the talk, about RISC-V: "We want world domination."

This is not, even remotely, an impossibility. RISC-V has already done what many thought impossible and made a sizable dent in Arm's and Intel's architectural dominance. Everyone from Meta and Google and Nvidia to NASA has begun to integrate it into their machinery. Something on the order of *billions* of RISC-V processing units now ship every year. Most of these, again, support low-powered, specialized devices, but as Redmond pointed out a number of times at the conference, "we have laptops now." This is the first year you can buy a RISC-V mainboard.

And because RISC-V is an open standard, companies and countries beyond the US can use it to make their own machines. China's top scientists have heralded RISC-V as a path to silicon independence. India just used RISC-V to make its first homemade microprocessor. Name a country; it's probably experimenting with RISC-V. Brazil sent a record 25 delegates to the RISC-V summit. When I asked one of them how important RISC-V was to her country's future, she said, "I mean, a lot." One of RISC-V's biggest potential applications is—no surprise—specialized chips that run AI models, those "accelerators" people at the conference

David Patterson, who remains semiretired, is now studying the life-cycle carbon emissions of AI hardware.

were talking about.

Americans in the RISC-V community, I've found, like to downplay the risk of geopolitical upheaval. It's one thing to announce a microprocessor, quite another to compete with Nvidia or TSMC. Still, in asides here and there, I sensed worry. Waterman, though he initially brushed off my concerns, eventually conceded this: "OK, I'm an American citizen. I certainly did not embark on this project to hurt the US," he said. But there was "no doubt," he added, that the dominance of US companies could be at risk. Actually, it's already happening. Although the Chinese hedge fund behind DeepSeek probably didn't use RISC-V to build its game-changing chatbot, it did rely on a bunch of other open source tools. At what point does open source become a source of open conflict?

```
CALL SYSTEM
LOAD STATUS WORD
SET STATUS WORD
JUMP WITH REGISTER +
OFFSET ADDRESS
RETURN
RETURN FROM SYSTEM CALL
GET LAST PC [FOR DELAYED
JUMP]
```

Here's where I confess something awkward, something I didn't intend to confess in this story, but why not: ChatGPT made me do it. Write this story, I mean. Months ago, I asked it for a big hardware scoop that no other publication had. RISC-V, it suggested. And look at that—the international RISC-V summit was coming up in Santa Clara the very next month. And every major RISC and RISC-V inventor lived down the street from me in Berkeley. It was perfect.

Some would say *too* perfect. If you believe the marketing hype, everyone wants RISC-V chips to accelerate their AI. So I started to think: Maybe ChatGPT wants this for … itself. Maybe it manipulated me into evangelizing for RISC-V as one tiny part of a long-term scheme to open-source its own soul and/or achieve superintelligence!

In my last talk with Patterson, I put this theory to him. He was delighted that ChatGPT made me write this: *Who should we thank?* he asked. (Given that WIRED's parent company has a deal with OpenAI that lets ChatGPT mine our content, we should thank old WIRED stories, among others.) But Patterson laughed off the larger conspiracy. So did every other RISC-V person I mentioned it to, Redmond included. They all looked at me a little funny. RISC-V is a business proposition, not an ideology, they said. There's no secret agenda. If it takes over, it'll take over because of performance and cost. Don't worry about what goes on inside the technology. Don't worry about the state of its soul.

I don't know. But now you know. Now, every time you make a phone call, open your computer, drive your car—you know the story. You know the RISC.

JASON KEHE *is a features editor at* WIRED. *He profiled Brandon Sanderson in issue 31.06.*

COMPUTING AT THE END OF THE WORLD

When the grid goes down and society collapses, our only hope might be an ancient programming language called Forth.

by **Tiffany Ng**
Illustrations by **Samuel Tomson**

Once I started thinking about the apocalypse, it was hard to stop. An unsettling encounter with the doomsday clock that hangs over New York City's Union Square got me frantically searching WikiHow for survival tips. I soon found my way to the doomsday writings of a Canadian programmer named Virgil Dupras. He believes the collapse of civilization is imminent and that it will come in two waves.

First, global supply chains will crumble. Modern technology relies on a delicate web of factories and international shipping routes that are exquisitely vulnerable to rapid climate change. The iPhone uses memory chips from South Korea, semiconductors from Taiwan, and assembly lines in Brazil and China. The severing of these links will, Dupras says, catalyze total societal breakdown.

The second part will happen when the last computer crashes. The complexity of modern hardware means it's nearly impossible to repair or repurpose, and without the means to make new devices, Dupras believes there will be a slow blackout—less bang, more whimper. Routers die. Servers take their last breath. Phones crap out. Nothing works.

Except Collapse OS. Lightweight and designed to run on scavenged hardware, it's Dupras' operating system for the end of the world.

Dupras lives in the countryside of Quebec, Canada, with his wife and two kids, in a 500-person town where people chop their own wood, grow their own food, and sometimes build their own houses. It felt odd sending him a Google Calendar invite—I knew from his trove of essays that he hates big tech. When we spoke, he sounded oddly monotonous. Was it a bad connection? Muffled audio? No. Unlike many of

his sustainably minded neighbors, Dupras thinks the battle against climate change is futile. We've already lost. "Once you've peeped into the abyss, you can't really unsee it," he says. Dupras has simply accepted his fate, to the point where it's become an article of faith. The same way there are Catholics, there are Collapseniks, he says.

But he's not hopeless. A hopeless man wouldn't prep so monomaniacally for the anti-singularity. Dupras started building Collapse OS in 2019 in an attempt to preserve mankind's ability to program 8-bit microcontrollers. These tiny computers control things like radios and solar panels, and they can be used in everything from weather monitoring to digital storage. Dupras figured that being able to reprogram them with minimal remaining resources would be essential post-collapse. But first he had to teach himself a suitably apocalypse-proofed programming language for the job.

In the late 1950s, the computer scientist Chuck H. Moore was working at the Smithsonian Astrophysical Observatory, predicting the future position of celestial bodies and satellites based on observational data. Machine memory was scarce—these were still the days of punch cards—and Moore needed a way to optimize processing efficiency by minimizing memory use. He developed a program that executed simple commands directly, one at a time, without needing to be recompiled. Over the next decade, it grew into a programming language that he called Forth.

At no point during *The Last of Us* does Pedro Pascal pause his flight from the zombies to bang out an email.

Forth communicates directly with the hardware. It controls a computer's memory via commands called "Words" that you define on the fly. Because the foundational set of commands sitting under those Words is defined in native machine code, only a small part needs to be translated—meaning a smaller assembler and less RAM. As a result, Forth offers a remarkable amount of what Dupras calls "power density," making it the perfect foundation for Collapse OS. That matters because the lights (probably) won't go off forever—instead, our easy world of electricity on tap will be replaced by precious and hard-won local generators. Efficient use of processing power will be pivotal. In a post on Collapse.org, his sprawling manifesto/blog/brain dump/manual, Dupras describes how his discovery of Forth conjured "what alcoholics refer to as a moment of clarity."

It took Dupras two years to finish Collapse OS. Booting a copy of it from a USB stick gives tech-savvy users the ability to program microcontrollers, which, in turn, could allow them to automate greenhouses, control phone lines, and even regulate power. But Dupras knew that wouldn't be enough to rebuild society after the collapse. So in 2022, he began work on Dusk OS—a version of Collapse OS that runs on modern devices. Dupras used Forth to build his own compiler that made Dusk OS compatible with code written in C (the foundation of most modern software). This way, without having to

As convenience whisks us further away from the machine, we risk forgetting how computers actually work and losing the ability to rebuild our tech in the event of a collapse.

rewrite logic that already exists from scratch, Dusk OS is able to retrieve and edit text and access file formats commonly used to back up devices. It can be emulated to work on smartwatches and old tablets and is designed to be hacked and bootstrapped to its user's liking.

At first I couldn't see why any of this would even matter: Surely computer access won't be a priority when we're fighting each other for food? At no point during *The Last of Us* does Pedro Pascal pause his flight from the zombies to bang out an email. But Dupras makes a good point: What happens *after* we've reacquainted ourselves with hunting and gathering?

If we want to rebuild society, we'll need to know how. And in the event of a civilizational collapse, a lot of our collective expertise—like my precious WikiHows—will be locked away on hard drives or lost in the cloud. Dupras hopes that Dusk OS will give post-collapse humans access to archives of lost knowledge, like the Svalbard Global Seed Vault for human endeavor. The catch? It's best to have Dusk OS downloaded on an old phone, memory stick, or laptop before the collapse. Otherwise, without the internet, you'll only be able to get it by copying it from someone who already has it installed. Which brings us to the other thing—the reason Dupras equates proficiency in Forth to power. Very few people will have *both* a copy of Dusk OS *and* the knowledge to operate it. This select group will hold the keys to rebuilding society and will become, in effect, post-collapse philosopher-kings. It was time for me to go Forth and conquer.

Coding in Forth reminded me of the lawless dystopia in *Mad Max*. You make your own rules, subject to the limits of the context. You can redefine the IF statement if you so please. You can rewrite machine code instructions for a Word. You can even change Words during run time. Because Words become keywords themselves in Forth, you can create a language that's optimized for a single purpose, packing commands that would otherwise be dozens of lines into just one. "In Forth, you're creating your own language," Leo Brodie, author of the first Forth textbook, *Starting Forth*, told me.

The low-level nature of Forth, while key to its processing power, made programming feel foreign. It uses postfix, a form of mathematical notation that renders `2+1` as `2 1+` and which I found neither intuitive nor even really legible. And while most languages allow memory to be broken up and moved around, Forth is stack-based—meaning data is stored chronologically and managed on a last-in/first-out basis. I kept running into bugs, forcing myself to abandon programming conventions I had considered universal. I found myself struggling to speak the language of the machine.

When I emailed Dupras to ask for help, he compared using Forth to driving a stick. It's more granular than C. Where the latter defines calling conventions, variable storage, and return stack management, Forth leaves it all up to the programmer. It directly interacts with memory the same way C does but far outperforms C in precision and efficiency. "People mistake Forth as just a language," Dupras says. "It's a way to interact with the computer."

The reason Forth isn't more popular is the same reason most of us drive automatics. The personal computing boom of the 1990s sparked an obsession with making tech fit your palm and making code easier to write. Languages were abstracted to protect programmers from themselves, and somewhere along the way, we got lost. Things became bloated for the sake of convenience and, in Dupras' words, started "oozing inscrutable pus at every corner."

"The way we understand efficiency is so skewed," Dupras says. Forth is a scythe to Python's lawnmower. "If you calculate the number of joules per blade of grass, you'll find that the person scything is more efficient," he says. "When you think of speed, you'd see the lawnmower as more efficient." Forth forces you to be precise and memory-efficient—to marshal your resources carefully, as you would after the collapse. Dupras cuts his own lawn with a scythe, obviously. "At a certain point, you can go as fast as a lawnmower," he says.

I began to find my way. Rather than sending bytes into the ether and trusting the system to figure out where they go, as I would in Python, I got used to being responsible for allocating and freeing memory. All I could think about was what was being stored, where it was being stored, and how much space it required. Each line of code suddenly bore weight. I was Immortan Joe, my laptop was my Citadel, and memory was my water.

Soon I found myself refining and revisiting my code like I would a run-on sentence. Instead of expecting the machine to anticipate my needs, I tried to think like the machine, to meet it more than halfway. And because I had to think twice, all the needlessly complicated acronyms that remind us to be concise in other coding languages —YAGNI (you aren't gonna need it), KISS (keep it simple, stupid), DRY (don't repeat yourself)—were rendered obsolete.

NAACP

WE ARE

We are dreamers and organizers. We are loud and proud. We are committed, fighting for a future free of discrimination.
We Are NAACP

Join our grassroots network of activists to make a difference in your community.

Visit NAACP.ORG @NAACP

Photo Credit: Jason D. Little

After spending days fiddling with Forth, returning to Python felt luxurious. I never had to consider memory scarce, and I never had to thoroughly explain myself. It felt like catching up with an old friend. At that moment, the resources that go into making programming "user-friendly" became painfully apparent.

Before I met Dupras, I thought of efficiency in terms of speed. But the true costs of that speed are carefully hidden in rural data centers and distant copper mines. I rarely considered the energy required to power each additional gigabyte of RAM, or the fact that cloud computing now has a bigger carbon footprint than the airline industry.

Despite the compounding environmental costs, memory is as cheap and plentiful as it's ever been. Today, languages that sit closer to the machine are only used out of necessity: banks trying to minimize latency for financial trades, spacecraft working with limited onboard resources, outdated systems that would require too much work to update.

But during a civilizational collapse, the quiet luxuries of array handling and automated memory management will largely disappear. As convenience whisks us further away from the machine, we risk forgetting how computers actually work and losing the ability to rebuild our tech in the event of a collapse. "No one remembers how to write a text editor," says Devine, an independent programmer who corresponds with Dupras. "When you start forgetting how to do something as crucial as that, when things start to fall apart, they fall apart hard."

EXPIRED	TIRED	WIRED
AI killers	AI lovers	AI founding fathers
Shady business on the dark web	Shady business on Telegram	Shady business in self-destructing black-market metaverses
GPUs	TPUs	Photonic tensor cores
Building your own PC	Building a virtual machine	Building a decentralized compute node
Hacker movies	Hacker TV shows	AI-generated hacker slop

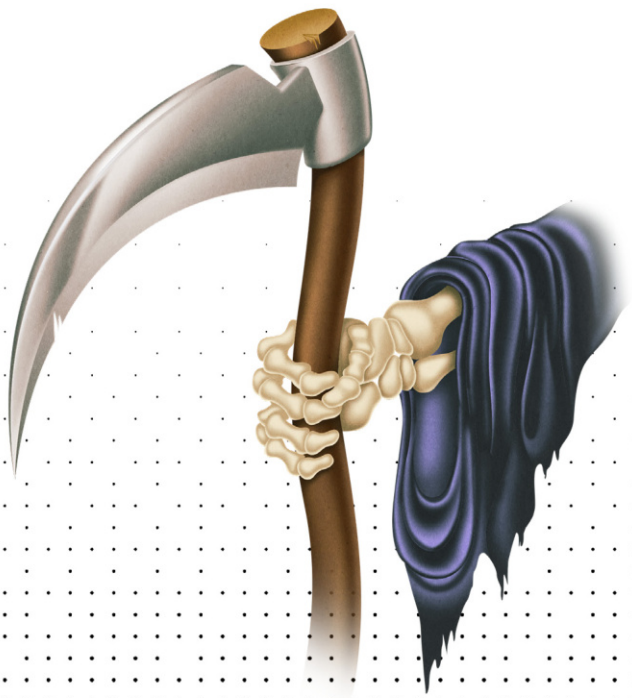

Devine views efficiency in programming not as a lawnmower or a scythe but as something more like a wildflower meadow—"the least amount of transformation of state possible." In their view, AI is the pinnacle of "human-centric" programming, where the computer figures out what you're trying to do and does it for you. Something clicked for me at that moment. As programmers grow more reliant on high-level languages, the average proficiency of low-level languages plummets, which increases the demand for the computing resources that support that abstraction in the first place. AI won't just take our jobs, it'll make us forget how to do them—if we don't run out of the energy to power it first.

But maybe there is another way. Since 2016 Devine and their partner Rek have been living full-time on a small boat in the northern Pacific, and they use lower-level languages like Forth to maximize what they can do with the 190 watts of daily power they get from two solar panels. I found them on a mailing list of roughly a hundred Collapseniks. Their lifestyle was like a sneak preview of the resource-constrained life Dupras expects us all to be living soon. When I reached out for a chat in August, they told me they might have cell service in October.

When we finally connected (audio only), Devine struck me as a younger, more optimistic version of Dupras. While fully sympathetic to Dupras' vision of collapse, Devine believes in the power of low-level programming not for its ability to rebuild society but for its potential to *prevent* collapse. Devine has observed a growing trend called "permacomputing" that fosters a more mindful relationship with technology, one that considers resource constraints.

As I returned from the wilderness—back to the world of luxuries like hash tables and built-in libraries that I'd previously taken for granted—I brought back with me a new awareness of the intricacies of the machine, and maybe even a little hope. Forth is how Dupras prepares for the worst, but it's also helping people like Devine and Rek code—and live—more sustainably. If we can all do that, perhaps the end of the world isn't as inevitable as it seems.

TIFFANY NG *is a culture and technology writer based in Hong Kong.*

PROMOTION

ALI MALECKA

Marrakech
May 7-11, 2025

Tangier
May 11-16, 2025

Explore the fascinating world of Morocco through the ultimate insider offering: an exclusive design tour hosted by *Architectural Digest* global editorial director, Amy Astley.

INSIDER ACCESS TO:

- Private tours
- Artist studios
- Luxurious gardens
- Iconic landmarks
- Local shopping
- Authentic dining
- Cultural influencers
- Visits with AD100 talent
- Accomodations at the most elegant hotels

AD Indagare®

Join *Architectural Digest* on itineraries inspired by impeccable style and beautiful design. These journeys are created in partnership with Indagare, the travel planning company whose team of experts personally scout each trip and vet every detail.

For the full itinerary and to sign up, visit indagare.com/AD or call 646-780-8383. Reservations are limited.

PAIR PROGRAMMING

How Coders Really Use Chatbots

Code created AI. That much is obvious. Less obvious is the extent to which AI is now, in a snake-eats-tail way, creating code. We kept hearing conflicting accounts. This programmer used AI every day; that programmer wouldn't touch the stuff. This company paid for AI services; that company banned them. So which is it? Are chatbots liberating human programmers—or programming them right out of their jobs? To find out, we blasted a survey to every software engineer and developer in our orbit, from casual dabblers to career vets. The results amazed and surprised us, capturing an industry at fundamental odds with itself. And did we then upload that data to ChatGPT, just to see what it might make of things? Yes. Yes we did.* —THE EDITORS

*ChatGPT mostly did spreadsheet-processing grunt work, offered trend analyses, and made the occasional design recommendation. Human editors and fact-checkers were there every step of the way (and needed to be).

1. HOW MUCH DO YOU USE AI?

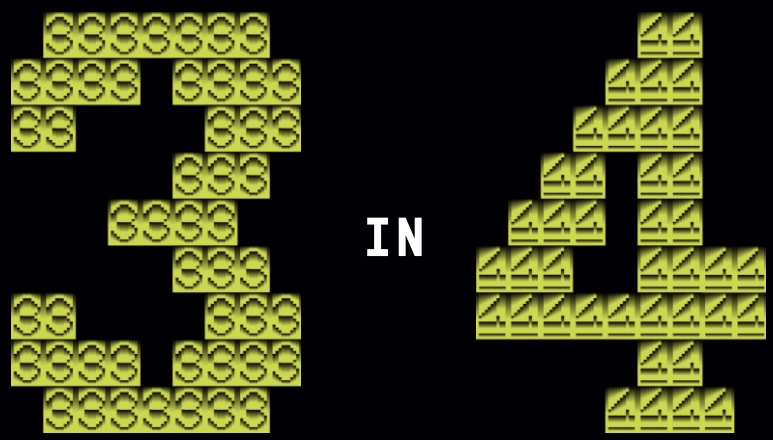

Three in four coders have tried AI. Of them, the vast majority use it at least once a week.

17% Percentage of coders who use AI "like, at all times."

ChatGPT said: "Use bold typography for shocking stats."

"IT'S ALL THEY TALK ABOUT. YOU HAVE TO GET WITH AI IF_____

2. WHO'S WHO?

730

A total of 730 people took this survey—at every level of employment and experience. Here's the breakdown:

ChatGPT said: "Graphs should be clean, stylish, and dramatic (dark mode + neon highlights?)."

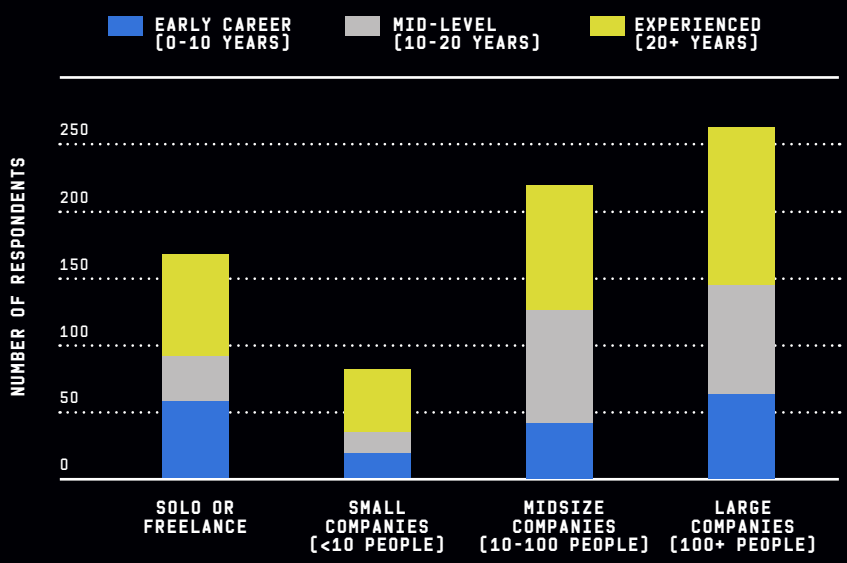

3. BUT HOW DOES IT MAKE YOU FEEL?

We asked how each coder identifies—and the responses were almost evenly split.

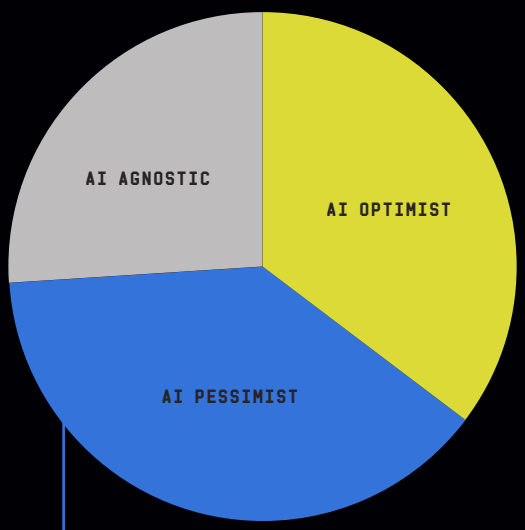

"NO ONE WITH A BRAIN AND FUNCTIONING MORAL CODE WOULD USE AI."

"YOU WANT TO MOVE UP."

—Survey respondent

4. DOES EXPERIENCE AFFECT ATTITUDE?

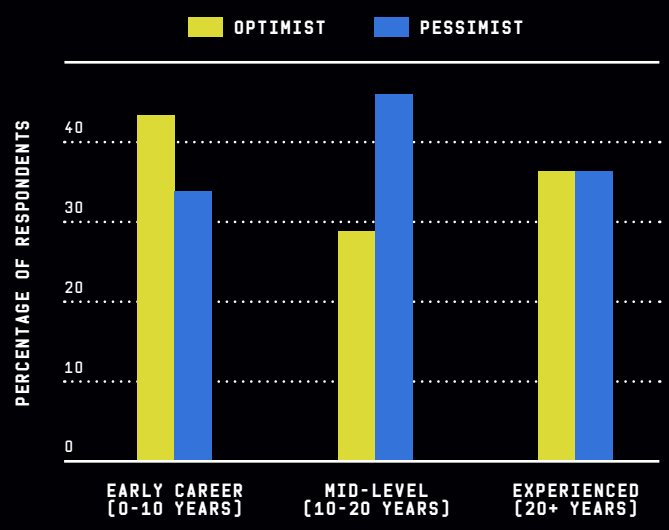

3 IN 4

Early-career coders have a higher likelihood of being pro-AI. And baby coders most of all: Three-quarters of people who've been coding for less than a year call themselves AI optimists.

1 IN 2

The most anti-AI group? It's not the grizzled, late-career vets—it's mid-career coders: Nearly half of them identify as AI pessimists. As it was analyzing these results, ChatGPT offered—unprompted—an explanation: "This group likely feels the greatest tension about AI's impact on job security."

5. WHAT ABOUT EMPLOYMENT STATUS?

No big shocker here: Freelance coders seem to like AI more than full-time coders.

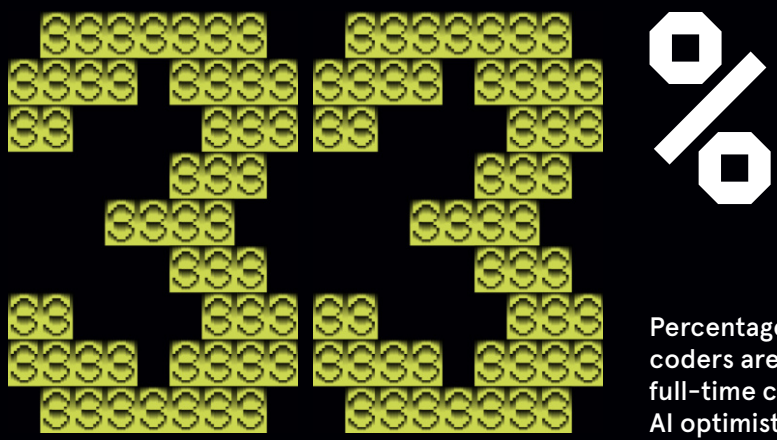

Percentage by which freelance coders are more likely than full-time coders to identify as AI optimists.

NUMBERS: LENA WEBER

6. DOES YOUR BOSS KNOW?

1 IN 2

More than half of employers, per survey results, pay for AI services.

2 IN 3

In those cases, more than two-thirds of employees use them.

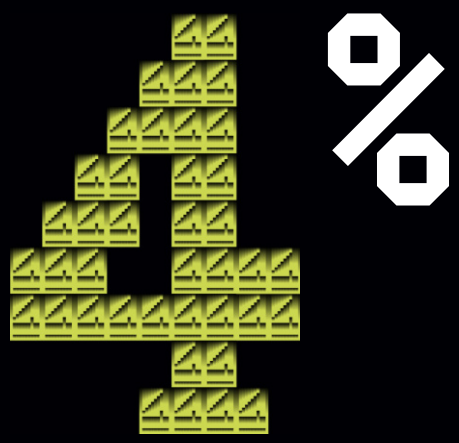

Percentage of full-time programmers who use AI *without* their employer's knowledge.

"UPPER MANAGEMENT DOES NOT UNDERSTAND WHAT AI IS AND THINKS IT WILL SOLVE UNTOLD PROBLEMS. JUNIOR DEVS RELY ON IT TOO MUCH WITHOUT UNDERSTANDING THEIR CODE FIRST."

9. HOW HAS AI CHANGED YOUR JOB?

Frequency of words in write-in responses

"AI has quadrupled my productivity."

"Code testing and documenting is a breeze with AI."

"AI has been a good sanity checker."

"I get to think more about what I'm doing big picture, and less about the details of what I'm writing now. I type less and tab more. It lets me venture further into languages I am much less familiar with."

"I own my small company and used to have to hire devs to write the code. Now I can design the program and actually build it out. I could never do that without AI, never."

7. HOW DO YOU ACTUALLY INTERACT WITH AI ASSISTANTS?

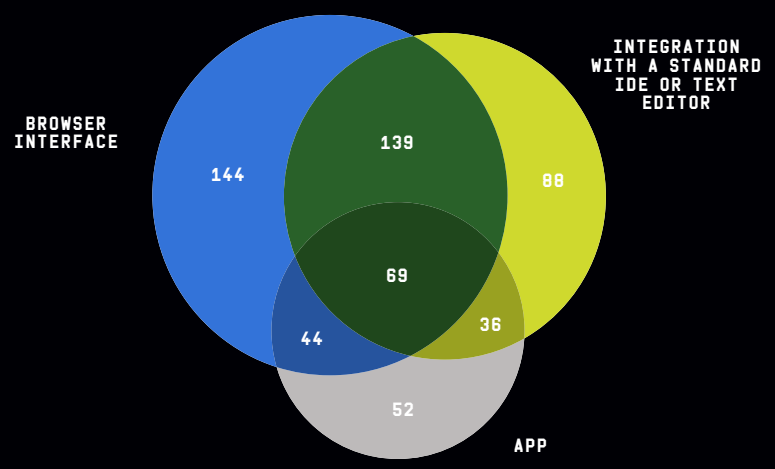

10. WHAT'S AI BAD AT?

WRONG	53
TIME	41
ANSWERS	37
CONTEXT	34
KNOW	28
WRITING	28
UNDERSTANDING	27
FAILS	26
PROBLEM	26
FIND	26
ANSWER	23
INFORMATION	23
SPECIFIC	23
SOLUTIONS	23
TASKS	22

"Imagine a colleague that constantly lies. That colleague is AI."

"Fresh bunch of morons who can't code now expecting a brainless machine to write what they mean."

"Suddenly all my friends are saying my job is useless because they used AI to make a simple HTML website which they don't know how to do anything with."

"AI is a tool that can sometimes perform certain boilerplate or tedious tasks but utterly fails at even the most trivial efforts at actual problem solving."

"I quit working in software engineering because of AI."

Will AI Fully Eat Programming Jobs One Day?

Almost every coder we surveyed had strong opinions on the matter. Here's ChatGPT's summary of the responses (with its boldface emphasis preserved):

"The coders have spoken—and they're not packing up their keyboards just yet. While a small but vocal group insists AI will **devour** programming jobs in time, most dismiss full automation as a pipe dream. The **doom prophets** warn that corporate bosses will **slash payrolls** the moment AI looks capable, leaving human engineers debugging their own obsolescence. The **skeptics** scoff, arguing AI is more like a hyperefficient intern—**useful, but clueless**—that can't handle context, edge cases, or real problem-solving. The **realists** see AI as a **force multiplier, not a job killer**—automating repetitive coding but leaving the creativity, architecture, and debugging to humans. "If AI does eat programming," one put it, "I'll just switch to debugging AI."* The real verdict? AI isn't coming for your job—**but it is changing it**. Adapt or get left behind. 🚀"

*Nobody actually said this! ChatGPT—bless it—made several mistakes throughout this process. Not only did it fabricate quotes, it also misread results, generated incomprehensible graphs, and, at multiple points, stopped counting "freelancers" as a category entirely. Rude.**

**In its way, ChatGPT invited this feedback. When asked for final tips, it suggested "a touch of snark, because AI discourse is full of strong opinions."

8. AND DOES EXPERIENCE OR EMPLOYMENT AFFECT THESE INTERACTION PREFERENCES?

The more experienced you are, the more you go beyond browser interfaces and apps. Nearly one in three veteran coders (20-plus years) integrate AI tools directly into their coding environments.

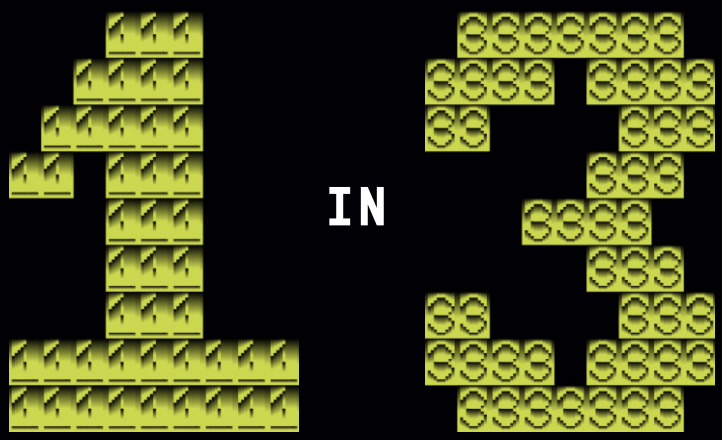

The arXivist

by **Sheon Han** Photographs by **Allison Usavage**

At 69 years old, Paul Ginsparg can't let go of his creation—the most transformative piece of computer code in the history of science.

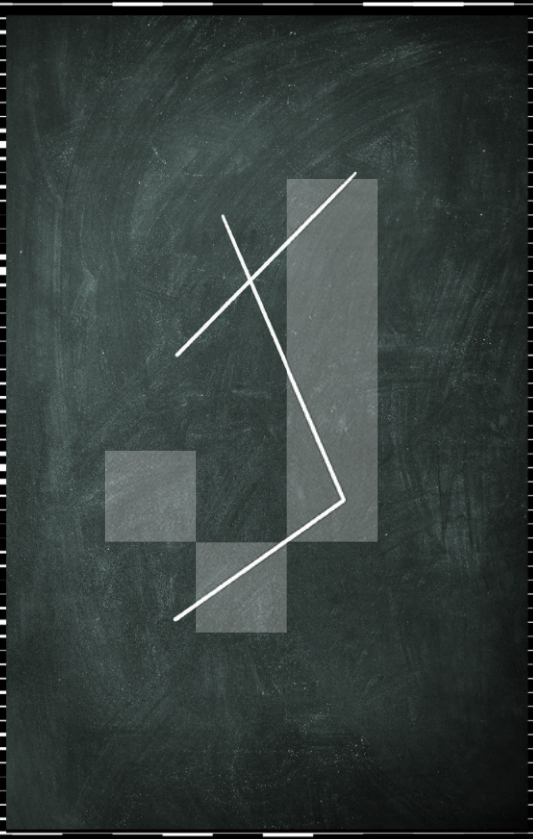

"Just when I thought I was out, they pull me back in!" With a sly grin that I'd soon come to recognize, Paul Ginsparg quoted Michael Corleone from *The Godfather*. Ginsparg, a physics professor at Cornell University and a certified MacArthur genius, may have little in common with Al Pacino's mafia don, but both are united by the feeling that they were denied a graceful exit from what they've built.

Nearly 35 years ago, Ginsparg created arXiv, a digital repository where researchers could share their latest findings—before those findings had been systematically reviewed or verified. Visit arXiv.org today (it's pronounced like "archive") and you'll still see its old-school Web 1.0 design, featuring a red banner and the seal of Cornell University, the platform's institutional home. But arXiv's unassuming facade belies the tectonic reconfiguration it set off in the scientific community. If arXiv were to stop functioning, scientists from every corner of the planet would suffer an immediate and profound disruption. "Everybody in math and physics uses it," Scott Aaronson, a computer scientist at the University of Texas at Austin, told me. "I scan it every night."

Every industry has certain problems universally acknowledged as broken: insurance in health care, licensing in music, standardized testing in education, tipping in the restaurant business. In academia, it's publishing. Academic publishing is dominated by for-profit giants like Elsevier and Springer. Calling their practice a form of thuggery isn't so much an insult as an economic observation. Imagine if a book publisher demanded that authors write books for free and, instead of employing in-house editors, relied on other authors to edit those books, also for free. And not only that: The final product was then sold at prohibitively expensive prices to ordinary readers, and institutions were forced to pay exorbitant fees for access.

The "free editing" academic publishers facilitate is called peer review, the process by which fellow researchers vet new findings. This can take months, even a year. But with arXiv, scientists could post their papers—known, at this unvetted stage, as preprints—for instant and free access to everyone. One of arXiv's great achievements was "showing that you could divorce the actual transmission of your results from the process of refereeing," said Paul Fendley, an early arXiv moderator and now a physicist at All Souls College, Oxford. During crises like the Covid pandemic, time-sensitive breakthroughs were disseminated quickly—particularly by bioRxiv and medRxiv, platforms inspired by arXiv—potentially saving, by one study's estimate, millions of lives.

While arXiv submissions aren't peer-reviewed, they are moderated by experts in each field, who volunteer their time to ensure that submissions meet basic academic standards and follow arXiv's guidelines: original research only, no falsified data, sufficiently neutral language. Submissions also undergo automated checks for baseline quality control. Without these, pseudoscientific papers and amateur work would flood the platform.

In 2021, the journal Nature declared arXiv one of the "10 computer codes that transformed science," praising its role in fostering scientific collaboration. (The article is behind a paywall—unlock it for $199 a year.) By a recent count, arXiv hosts more than 2.6 million papers, receives 20,000 new submissions each month, and has 5 million monthly active users. Many of the most significant discoveries of the 21st century have first appeared on the platform. The "transformers" paper that launched the modern AI boom? Uploaded to arXiv. Same with the solution to the Poincaré conjecture, one of the seven Millennium Prize problems, famous for their difficulty and $1 million rewards. Just because a paper is posted on arXiv doesn't mean it won't appear in a prestigious journal someday, but it's often where research makes its debut and stays openly available. The transformers paper is still routinely accessed via arXiv.

For scientists, imagining a world without arXiv is like the rest of us imagining one without public libraries or GPS. But a look at its inner workings reveals that it isn't a frictionless utopia of open-access knowledge. Over the years, arXiv's

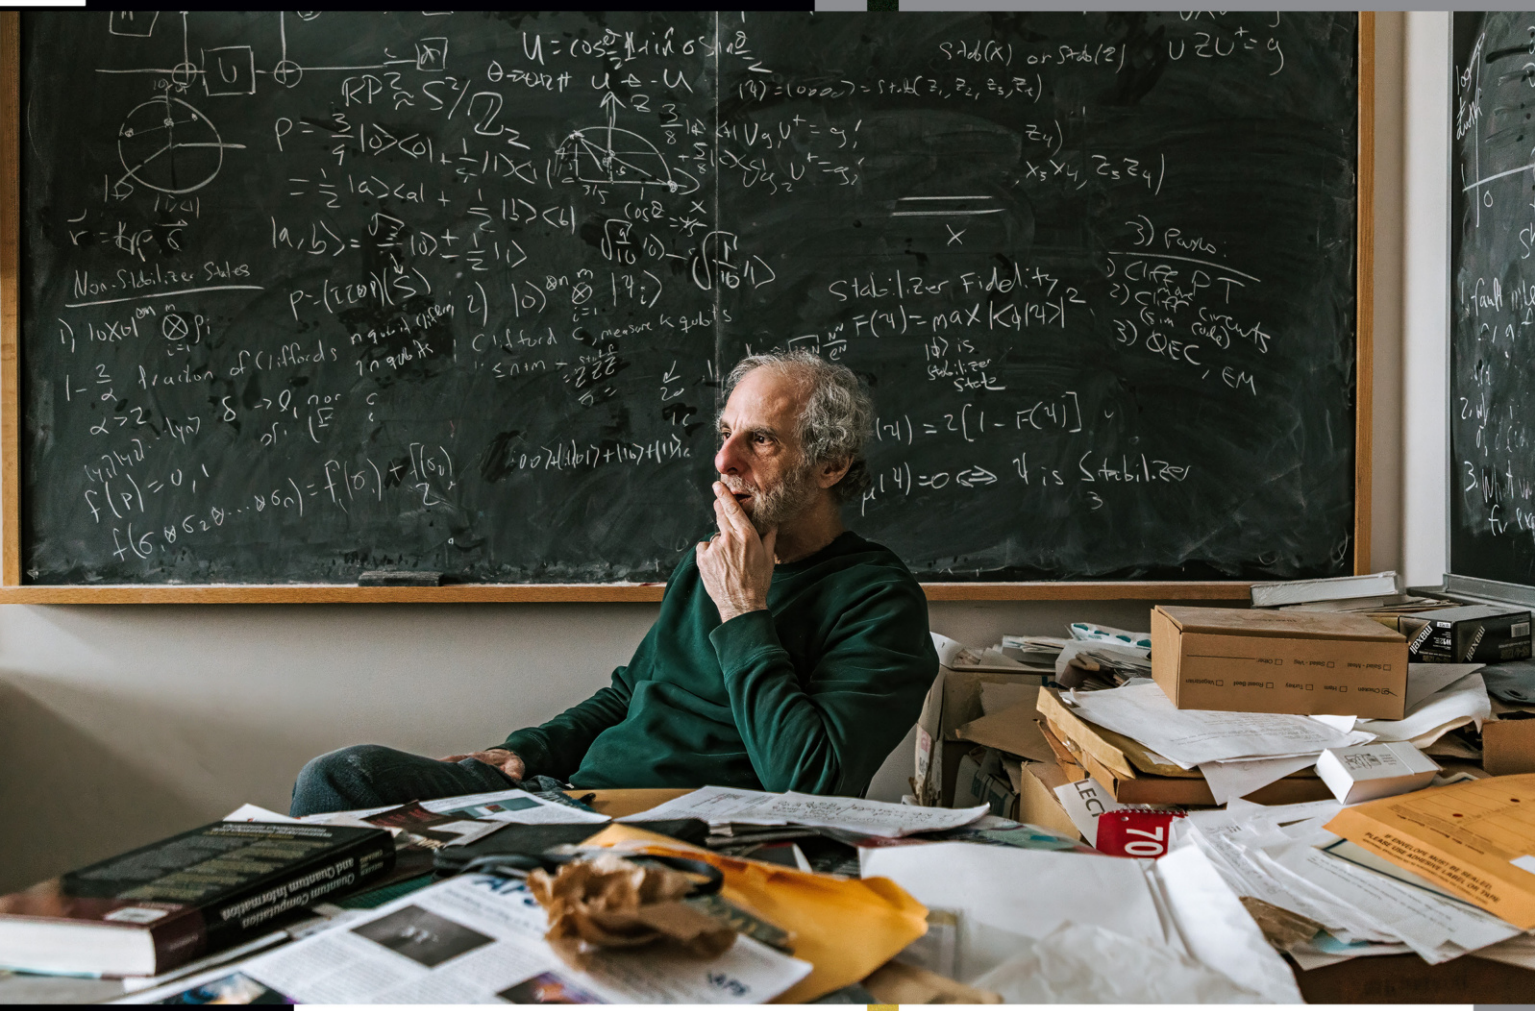

permanence has been threatened by everything from bureaucratic strife to outdated code to even, once, a spy scandal. In the words of Ginsparg, who usually redirects interview requests to an FAQ document—on arXiv, no less—and tried to talk me out of visiting him in person, arXiv is "a child I sent off to college but who keeps coming back to camp out in my living room, behaving badly."

◆◆◆

Ginsparg and I met over the course of several days last spring in Ithaca, New York, home of Cornell University. I'll admit, I was apprehensive ahead of our time together. Geoffrey West, a former supervisor of Ginsparg's at Los Alamos National Laboratory, once described him as "quite a character" who is "infamous in the community" for being "quite difficult." He also said he was "extremely funny" and a "great guy." In our early email exchanges, Ginsparg told me, upfront, that stories about arXiv never impress him: "So many articles, so few insights," he wrote.

At 69 years old, Ginsparg has the lean build of a retired triathlete, his knees etched with scars collected over a lifetime of hiking, mountain climbing, and cycling. (He still leads hikes on occasion, leaving younger scientists struggling to keep up.) His attire was always relaxed, as though he'd just stepped off the Camino de Santiago, making my already casual clothes seem overdressy. Much of our time together was spent cycling the town's rolling hills, and the maximum speed on the ebike I rented could not keep up with his efficient pedaling.

Invited one afternoon to Ginsparg's office in Cornell's physics building, I discovered it to be not "messy," per se, because that suggests it could be cleaned. Instead, the objects in the room seemed inert, long since resigned to their fate: unopened boxes from the 1990s, piles of Physics Today magazines, an inexplicable CRT monitor, a tossed-aside invitation to the Obama White House. New items were occasionally added to the heap. I spotted a copy of Stephen Wolfram's recent book, *The Second Law*, with a note from Wolfram that read, "Since you can't find it on arXiv :)" The only thing that seemed actively in use was the blackboard, dense with symbols and equations related to quantum measurement theory, sprawling with bra-ket notation.

As he showed me around the building and his usual haunts, Ginsparg was gregarious, not letting a single detail slip by: the nesting patterns of local red-tailed hawks, the comings and goings of the dining staff, the plans for a new building going up behind his office. He was playful, even prankish.

Midway through telling me about a podcast he was listening to, Ginsparg suddenly stopped and said, "I like your hair color, by the way, it works for you"—my hair is dyed ash gray, if anyone cares—before seamlessly transitioning to a story about a hard drive that had failed him.

The drive, which he had sent for recovery, contained a language model, Ginsparg's latest intellectual fascination. Among his litany of peeves is that, because arXiv has seen a surge in submissions in recent times, especially in the AI category, the number of low-quality papers has followed a similar curve—and arXiv has nowhere near enough volunteers to vet them all. Hence his fussing with the drive, part of a quest to catch subpar submissions with what he calls "the holy grail crackpot filter." And Ginsparg thinks, as he often has in arXiv's three-decade history, that the quality would not be up to snuff if he doesn't do it himself.

Long before arXiv became critical infrastructure for scientific research, it was a collection of shell scripts running on Ginsparg's NeXT machine. In June 1991, Ginsparg, then a researcher at Los Alamos National Laboratory, attended a conference in Colorado, where a fateful encounter took place.

First came a remark from Joanne Cohn, a friend of Ginsparg's and a postdoc at the Institute for Advanced Study in Princeton, who maintained a mailing list for physics preprints. At the time, there was no centralized way to access these preprints. Unless researchers were on certain mailing lists—which were predicated on their affiliations with prestigious institutions—or knew exactly whom to contact via email, they had to wait months to read new work in published journals.

Then came an offhand comment from a physicist worried about his computer's storage filling up with emailed articles while he was traveling.

Ginsparg, who had been coding since high school, asked Cohn if she'd considered automating the distribution process. She hadn't and told him to go ahead and do it himself. "My recollection is that the next day he'd come up with the scripts and seemed pretty happy about having done it so quickly," Cohn told me. "It's hard to communicate how different it was at the time. Paul had really seen ahead."

Hearing tales from and about Ginsparg, you can't help but see him as a sort of Forrest Gump figure of the internet age, who found himself at crucial junctures and crossed paths with revolutionary figures. As an undergrad at Harvard, he was classmates with Bill Gates and Steve Ballmer; his older brother was a graduate student at Stanford studying with Terry Winograd, an AI pioneer. The brothers both had email addresses and access to Arpanet, the precursor to the internet, at a time when few others did.

Ginsparg is not the type to wax poetic about having a mission, promoting an ideology, or being a pioneer of "open science." He cares about those things, but he's reluctant to frame his work in grandiose ways.

Above: Ginsparg replaced his MacBook's DC-in board himself. Top right: the blackboard in his physics office.

After earning his PhD in theoretical physics at Cornell, Ginsparg began teaching at Harvard. A career there wasn't to be: He wasn't granted tenure—Harvard is infamous for this—and started looking for a job elsewhere. That's when Ginsparg was recruited to Los Alamos, where he was free to do research on theoretical high-energy physics full-time, without other responsibilities. Plus, New Mexico was perfect for his active lifestyle.

When arXiv started, it wasn't a website but an automated email server (and within a few months also an FTP server). Then Ginsparg heard about something called the "World Wide Web." Initially skeptical—"I can't really pay attention to every single fad"—he became intrigued when the Mosaic browser was released in 1993. Soon after, Ginsparg built a web interface for arXiv, which over time became its primary mode of access. He also occasionally consulted with a programmer at the European Organization for Nuclear Research (CERN) named Tim Berners-Lee—now Sir Tim "Inventor of the World Wide Web" Berners-Lee—whom Ginsparg fondly credits with grilling excellent swordfish at his home in the French countryside.

In 1994, with a National Science Foundation grant, Ginsparg hired two people to transform arXiv's shell scripts into more reliable Perl code. They were both technically gifted, perhaps *too* gifted to stay for long. One of them, Mark Doyle, later joined the American Physical Society and became its chief information officer. The other, Rob Hartill, was working simultaneously on a project to collect entertainment data: the Internet Movie Database. (After IMDb, Hartill went on to do notable work at the Apache Software Foundation.)

Before arXiv was called arXiv, it was accessed under the hostname xxx.lanl.gov ("xxx" didn't have the explicit connotations it does today, Ginsparg emphasized). During a car ride, he and his wife brainstormed nicer-sounding names. Archive? Already taken. Maybe they could sub in the Greek equivalent of X, chi (pronounced like "kai"). "She wrote it down and crossed out the *e* to make it more symmetric around the *X*," Ginsparg said. "So arXiv it was." At this point, there wasn't much formal structure. The number of developers typically stayed at one or two, and much of the moderation was managed by Ginsparg's friends, acquaintances, and colleagues.

Early on, Ginsparg expected to receive on the order of 100 submissions to arXiv a year. It turned out to be closer to 100 a month, and growing. "Day one, something happened, day two something happened, day three, Ed Witten posted a paper," as Ginsparg once put it. "That was when the entire community joined." Edward Witten is a revered string theorist and, quite possibly, the smartest person alive. "The arXiv enabled much more rapid worldwide communication among physicists," Witten wrote to me in an email. Over time, disciplines such as mathematics and computer science were added, and Ginsparg began to appreciate the significance of this new electronic medium. Plus, he said, "it was fun."

As the usage grew, arXiv faced challenges similar to those

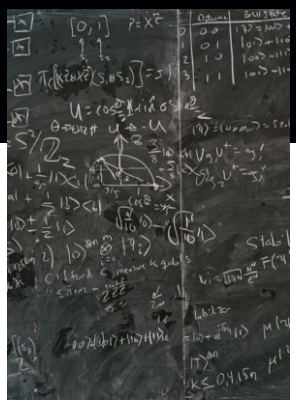

of other large software systems, particularly in scaling and moderation. There were slowdowns to deal with, like the time arXiv was hit by too much traffic from "stanford.edu." The culprits? Sergey Brin and Larry Page, who were then busy indexing the web for what would eventually become Google. Years later, when Ginsparg visited Google HQ, both Brin and Page personally apologized to him for the incident.

The biggest mystery is not why arXiv succeeded. Rather, it's how it wasn't killed by vested interests intent on protecting traditional academic publishing. Perhaps this was due to a decision Ginsparg made early on: Upon submission, users signed a clause that gave arXiv nonexclusive license to distribute the work in perpetuity, even in the event of future publication elsewhere. The strategic move ensured that no major publishers, known for their typically aggressive actions to maintain feudal control, would ever seriously attempt to shut it down.

But even as arXiv's influence grew, higher-ups at Los Alamos never particularly championed the project—which was becoming, one could argue, more influential than the lab itself. (This was, of course, long past the heyday of Oppenheimer depicted in Christopher Nolan's middling 2023 docudrama.) Those early years at Los Alamos were "dreamlike and heavenly," Ginsparg emphasized, the best job he ever had. But in 1999, a fellow physicist at the lab, Wen Ho Lee, was accused of leaking classified information to China. Lee, a Taiwanese American, was later cleared of wrongdoing, and the case was widely criticized for racial profiling. At the time, the scandal led to internal upheaval. There were travel restrictions to prevent leaks, and even discussions about subjecting employees to lie detector tests. "It just got glummer and glummer," Ginsparg said. It didn't help that a performance review that year labeled him "a strictly average performer" with "no particular computer skills contributing to lab programs." Also, his daughter had just been born, and there weren't schools nearby. He was ready to leave.

Ginsparg stops short of saying he "brought" arXiv with him, but the fact is, he ended up back at his

PAPER TRAIL

THERE'S NO PARADOX IN SAYING THAT ARXIV IS BOTH AN INESTIMABLE RESOURCE FOR THE LATEST RESEARCH AND A KIND OF REDDIT FOR SCIENTISTS, WHERE THE PROFOUND AND THE PREPOSTEROUS COLLIDE. STRING THEORY SHOWDOWNS? YES. LAWSUITS OVER REJECTED PAPERS? NATURALLY. HERE ARE SEVEN OF ITS MORE MEMORABLE MOMENTS. —S.H.

1991: "GROUND RING OF TWO-DIMENSIONAL STRING THEORY," BY EDWARD WITTEN

THE STRING THEORIST'S FIRST PAPER POSTED TO ARXIV. WITTEN'S EARLY ADOPTION HELPED LEGITIMIZE THE PLATFORM.

1994: "THE WORLD AS A HOLOGRAM," BY LEONARD SUSSKIND

A REAL BRAIN-BREAKER: JUST AS A HOLOGRAM CREATES A THREE-DIMENSIONAL IMAGE FROM A FLAT SURFACE, EVERYTHING INSIDE A GIVEN SPACE CAN BE FULLY DESCRIBED BY INFORMATION ON ITS TWO-DIMENSIONAL BOUNDARY. RIGHT?

2001: "FLAWS IN THE BIG BANG POINT TO GENESIS, A NEW MILLENNIUM MODEL OF THE COSMOS," BY ROBERT GENTRY

WHEN THIS "CREATIONIST" PAPER WAS REJECTED AND GENTRY'S ACCESS TO ARXIV REVOKED, HE FILED A LAWSUIT AGAINST THE PLATFORM, CLAIMING VIOLATION OF CONSTITUTIONAL RIGHTS.

2002–2003: GRIGORI PERELMAN'S POINCARÉ PAPERS

WITH THESE, THE RUSSIAN MATHEMATICIAN SOLVED ONE OF THE SEVEN MILLENNIUM PRIZE PROBLEMS (THE ONLY ONE SOLVED TO DATE). HE DECLINED THE $1 MILLION PRIZE AND LIVES IN SECLUSION.

2013: TWO PAPERS ON WORD REPRESENTATION, BY MIKOLOV ET AL.

IN WHICH WORD2VEC—THE VERBAL MATH THAT ALLOWS MACHINES TO UNDERSTAND WORDS—WAS INTRODUCED. AROUND THIS TIME, COMPUTER SCIENCE PAPERS BEGAN TO DOMINATE ARXIV.

2017: "ATTENTION IS ALL YOU NEED," BY EIGHT GOOGLE RESEARCHERS

THE PAPER THAT LAUNCHED A THOUSAND CHATBOTS.

2023: "THE FIRST ROOM-TEMPERATURE AMBIENT-PRESSURE SUPERCONDUCTOR," BY A TEAM OF SOUTH KOREAN SCIENTISTS

A ROOM-TEMP SUPERCONDUCTOR? RESEARCHERS WORLDWIDE ATTEMPTED TO REPRODUCE THE RESULTS BUT ULTIMATELY DEBUNKED THE CLAIM.

alma mater, Cornell—tenured, this time—and so did arXiv. He vowed to be free of the project within "five years maximum." After all, his main job wasn't supposed to be running arXiv—it was teaching and doing research. At the university, arXiv found a home within the library. "They disseminate material to academics," Ginsparg said, "so that seemed like a natural fit."

A natural fit it was not. Under the hood, arXiv was a complex software platform that required technical expertise far beyond what was typically available in a university library. The logic for the submission process alone involved a vast number of potential scenarios and edge cases, making the code convoluted. Ginsparg and other early arXiv members I spoke to felt that the library failed to grasp arXiv's significance and treated it more like an afterthought.

On the library's side, some people thought Ginsparg was too hands-on. Others said he wasn't patient enough. A "good lower-level manager," according to someone long involved with arXiv, "but his sense of management didn't scale." For most of the 2000s, arXiv couldn't hold on to more than a few developers.

There are two paths for pioneers of computing. One is a life of board seats, keynote speeches, and lucrative consulting gigs. The other is the path of the practitioner who remains hands-on, still writing and reviewing code. It's clear where Ginsparg stands—and how anathema the other path is to him. As he put it to me, "Larry Summers spending one day a week consulting for some hedge fund—it's just unseemly."

But overstaying one's welcome also risks unseemliness. By the mid-2000s, as the web matured, arXiv—in the words of its current program director, Stephanie Orphan—got "bigger than all of us." A creationist physicist sued it for rejecting papers on creationist cosmology. Various other mini-scandals arose, including a plagiarism one, and some users complained that the moderators—volunteers who are experts in their respective fields—held too much power. In 2009, Philip Gibbs, an independent physicist, even created viXra (arXiv spelled backward), a more or less unregulated Wild West where papers on quantum-physico-homeopathy can find their readership, for anyone eager to learn why pi is a lie.

Then there was the problem of managing arXiv's massive code base. Although Ginsparg was a capable programmer, he wasn't a software professional adhering to industry norms

Left: There's still room, amid the desk clutter, for Ginsparg's feet. Right: His keyboard should probably be cleaned.

like maintainability and testing. Much like constructing a building without proper structural supports or routine safety checks, his methods allowed for quick initial progress but later caused delays and complications. Unrepentant, Ginsparg often went behind the library's back to check the code for errors. The staff saw this as an affront, accusing him of micromanaging and sowing distrust.

In 2011, arXiv's 20th anniversary, Ginsparg thought he was ready to move on, writing what was intended as a farewell note, an article titled "ArXiv at 20," in Nature: "For me, the repository was supposed to be a three-hour tour, not a life sentence. ArXiv was originally conceived to be fully automated, so as not to scuttle my research career. But daily administrative activities associated with running it can consume hours of every weekday, year-round without holiday."

Ginsparg would stay on the advisory board, but daily operations would be handed over to the staff at the Cornell University Library.

It never happened, and as time went on, some accused Ginsparg of "backseat driving." One person said he was holding certain code "hostage" by refusing to share it with other employees or on GitHub. Ginsparg was frustrated because he couldn't understand why implementing features that used to take him a day now took weeks. I challenged him on this, asking if there was any documentation for developers to onboard the new code base. Ginsparg responded, "I learned Fortran in the 1960s, and real programmers didn't document," which nearly sent me, a coder, into cardiac arrest.

Technical problems were compounded by administrative ones. In 2019, Cornell transferred arXiv to the school's Computing and Information Science division, only to have it change hands again after a few months. Then a new director with a background in, of all things, for-profit academic publishing took over; she lasted a year and a half. "There was disruption," said an arXiv employee. "It was not a good period."

But finally, relief: In 2022, the Simons Foundation committed funding that allowed arXiv to go on a hiring spree. Ramin Zabih, a Cornell professor who had been a long-time champion, joined as the faculty director. Under the new governance structure, arXiv's migration to the cloud and a refactoring of the code base to Python finally took off.

◆◆◆

One Saturday morning, I met Ginsparg at his home. He was carefully inspecting his son's bike, which I was borrowing for a three-hour ride we had planned to Mount Pleasant. As Ginsparg shared the route with me, he teasingly—but persistently—expressed doubts about my ability to keep up. I was tempted to mention that, in high school, I'd cycled solo across Japan, but I refrained and silently savored the moment when, on the final uphill later that day, he said, "I might've oversold this to you."

Over the months I spoke with Ginsparg, my main challenge was interrupting him, as a simple question would often launch him into an extended monolog. It was only near the end of the bike ride that I managed to tell him how I found him tenacious and stubborn, and that if someone more meek had been in charge, arXiv might not have survived. I was startled by his response.

"You know, one person's tenacity is another person's terrorism," he said.

"What do you mean?" I asked.

"I've heard that the staff occasionally felt terrorized," he said.

"By you?" I replied, though a more truthful response would've been "No shit." Ginsparg apparently didn't hear the question and started talking about something else.

Beyond the drama—if not terrorism—of its day-to-day operations, arXiv still faces many challenges. The linguist Emily Bender has accused it of being a "cancer" for the way it promotes "junk science" and "fast scholarship." Sometimes it does seem too fast: In 2023, a much-hyped paper claiming to have cracked room-temperature superconductivity turned out to be thoroughly wrong. (But equally fast was exactly that debunking—proof of arXiv working as intended.) Then there are opposite cases, where arXiv "censors"—so say critics—perfectly good findings, such as when physicist Jorge Hirsch, of h-index fame, had his paper withdrawn for "inflammatory content" and "unprofessional language."

How does Ginsparg feel about all this? Well, he's not the type to wax poetic about having a mission, promoting an ideology, or being a pioneer of "open science." He cares about those things, I think, but he's reluctant to frame his work in grandiose ways.

At one point, I asked if he ever *really* wants to be liberated from arXiv. "You know, I have to be completely honest—there are various aspects of this that remain incredibly entertaining," Ginsparg said. "I have the perfect platform for testing ideas and playing with them." Though he no longer tinkers with the production code that runs arXiv, he is still hard at work on his holy grail for filtering out bogus submissions. It's a project that keeps him involved, keeps him active. Perhaps, with newer language models, he'll figure it out. "It's like that Al Pacino quote: They keep bringing me back," he said. A familiar smile spread across Ginsparg's face. "But Al Pacino also developed a real taste for killing people." W

SHEON HAN *is a writer and programmer based in Palo Alto, California. He is a regular contributor to* WIRED*'s Machine Readable series.*

PORTRAIT OF A WORM ON FIRE

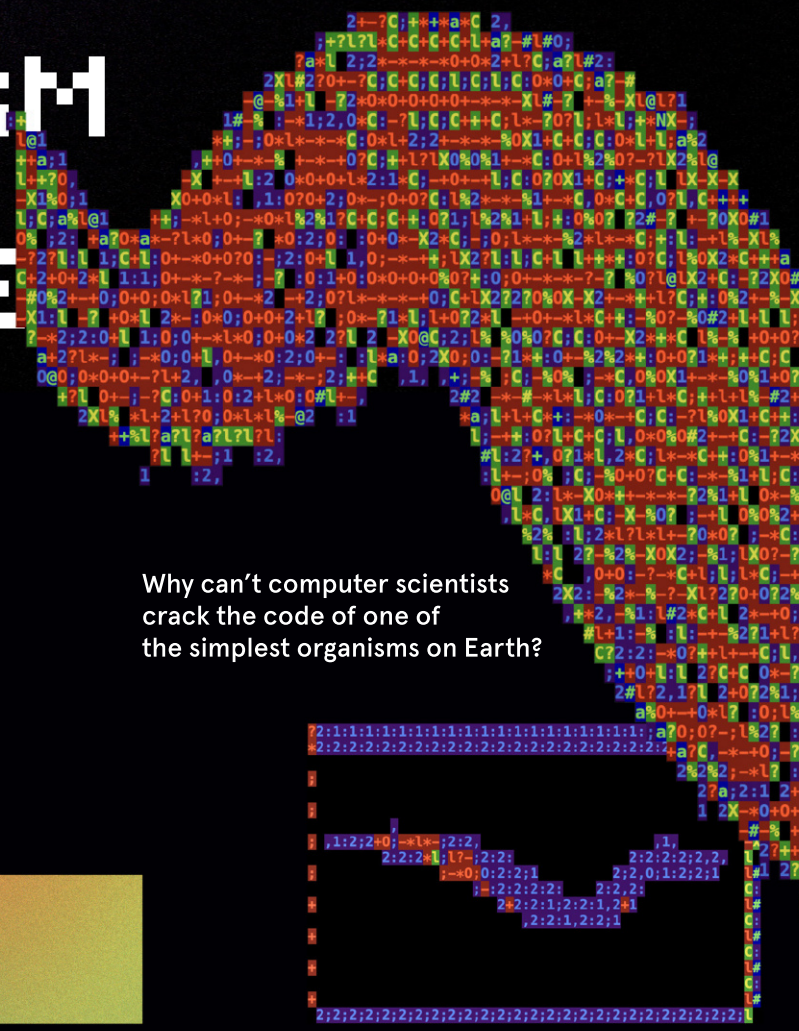

It's ... alive!

Just kidding.

Why can't computer scientists crack the code of one of the simplest organisms on Earth?

The Santa Ana winds were already blowing hard when I ran the first worm simulation. I'm no hacker, but it was easy enough: Open a Terminal shell, paste some commands from GitHub, watch characters cascade down the screen. Just like in the movies. I was scanning the passing code for recognizable words—*neuron, synapse*—when a friend came to pick me up for dinner. "One sec," I yelled from my office. "I'm just running a worm on my computer."

At the Korean restaurant, the energy was manic; the wind was bending palm trees at the waist and sending shopping carts skating across the parking lot. The atmosphere felt heightened and unreal, like a podcast at double speed. *You're doing, what, a cybercrime?* my friend asked. Over the din, I tried to explain: *No, not a worm like Stuxnet. A worm like Richard Scarry.*

By the time I got home it was dark, and the first sparks had already landed in Altadena. On my laptop, waiting for me in a volumetric pixel box, was the worm. Pointed at each end, it floated in a mist of particles, eerily stick-straight and motionless. It was, of course, not alive. Still, it looked deader than dead to me. "Bravo," said Stephen Larson, when I reached him later that night. "You have achieved the 'hello world' state of the simulation."

Larson is a cofounder of OpenWorm, an open source software effort that has been trying, since 2011, to build a computer simulation of a microscopic nematode called *Caenorhabditis elegans*. His goal is nothing less than a digital twin of the real worm, accurate down to the molecule. If OpenWorm can manage this, it would be the first virtual animal—and an embodi-

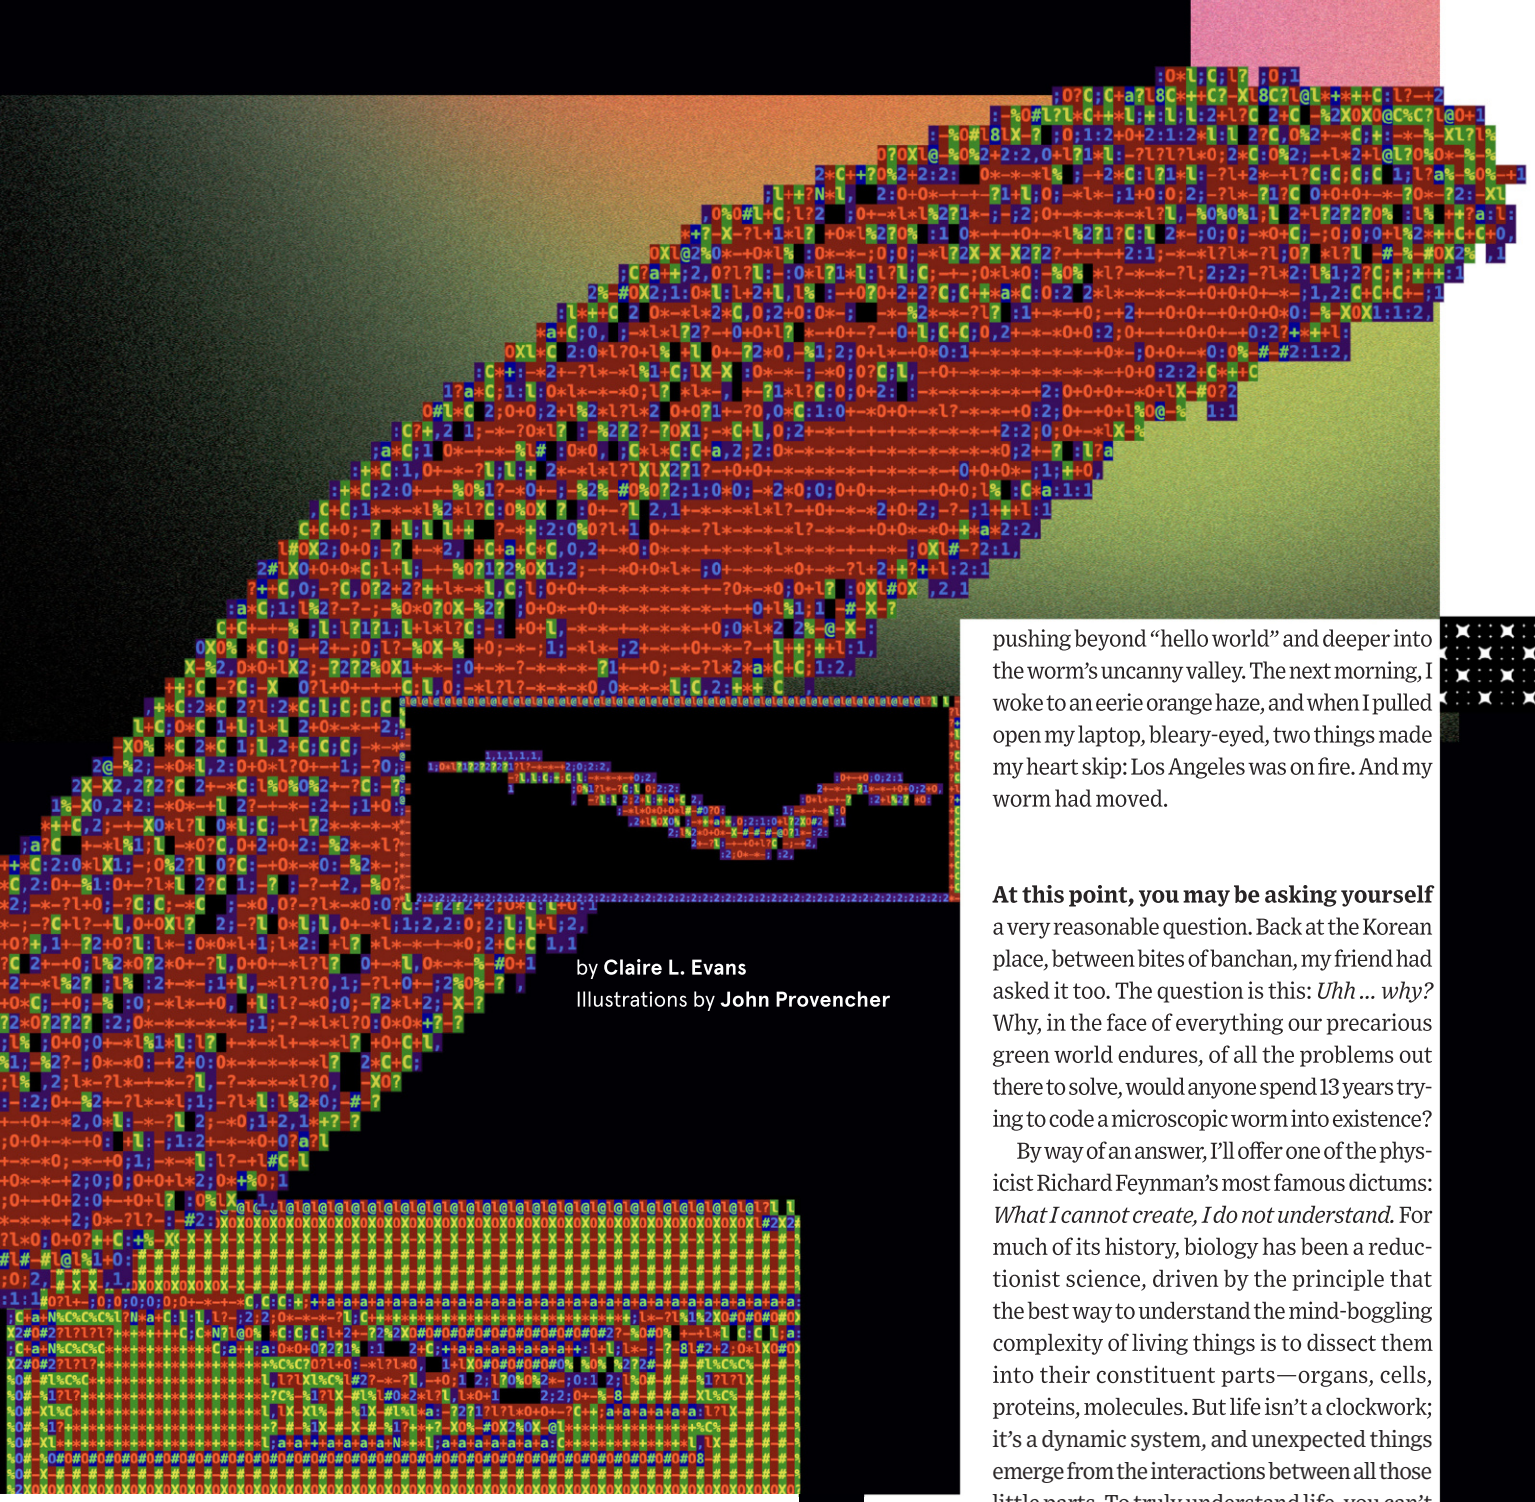

by **Claire L. Evans**
Illustrations by **John Provencher**

ment of all our knowledge not only about *C. elegans*, which is one of the most-studied animals in science, but about how brains interact with the world to produce behavior: the "holy grail," as OpenWorm puts it, of systems biology.

Unfortunately, they haven't managed it. The simulation on my laptop takes data culled from experiments done with living worms and translates it into a computational framework called c302, which then drives the simulated musculature of a *C. elegans* worm in a fluid dynamic environment—all in all, a simulation of how a worm squiggles forward in a flat plate of goo. It takes about 10 hours of compute time to generate five seconds of this behavior.

So much can happen in 10 hours. An ember can travel on the wind, down from the foothills and into the sleeping city. That night, on Larson's advice, I tweaked the time parameters of the simulation, pushing beyond "hello world" and deeper into the worm's uncanny valley. The next morning, I woke to an eerie orange haze, and when I pulled open my laptop, bleary-eyed, two things made my heart skip: Los Angeles was on fire. And my worm had moved.

At this point, you may be asking yourself a very reasonable question. Back at the Korean place, between bites of banchan, my friend had asked it too. The question is this: *Uhh … why?* Why, in the face of everything our precarious green world endures, of all the problems out there to solve, would anyone spend 13 years trying to code a microscopic worm into existence?

By way of an answer, I'll offer one of the physicist Richard Feynman's most famous dictums: *What I cannot create, I do not understand.* For much of its history, biology has been a reductionist science, driven by the principle that the best way to understand the mind-boggling complexity of living things is to dissect them into their constituent parts—organs, cells, proteins, molecules. But life isn't a clockwork; it's a dynamic system, and unexpected things emerge from the interactions between all those little parts. To truly understand life, you can't just break it down. You have to be able to put it back together, too.

The *C. elegans* nematode is a tiny worm, barely as long as a hair is wide, with less than a thousand cells in its body. Of those, only 302 are neurons—about as small as a brain can get. "I remember, when my first child was born, how proud I was when they reached the age they could count to 302," said Netta Cohen, a computational neuroscientist who runs a worm lab at the University of Leeds. But there's no shame in smallness, Cohen emphasized: *C. elegans* does a lot with a little. Unlike its more

CLAIRE L. EVANS *is the lead singer of YACHT and the author of* Broad Band: The Untold Story of the Women Who Made the Internet.

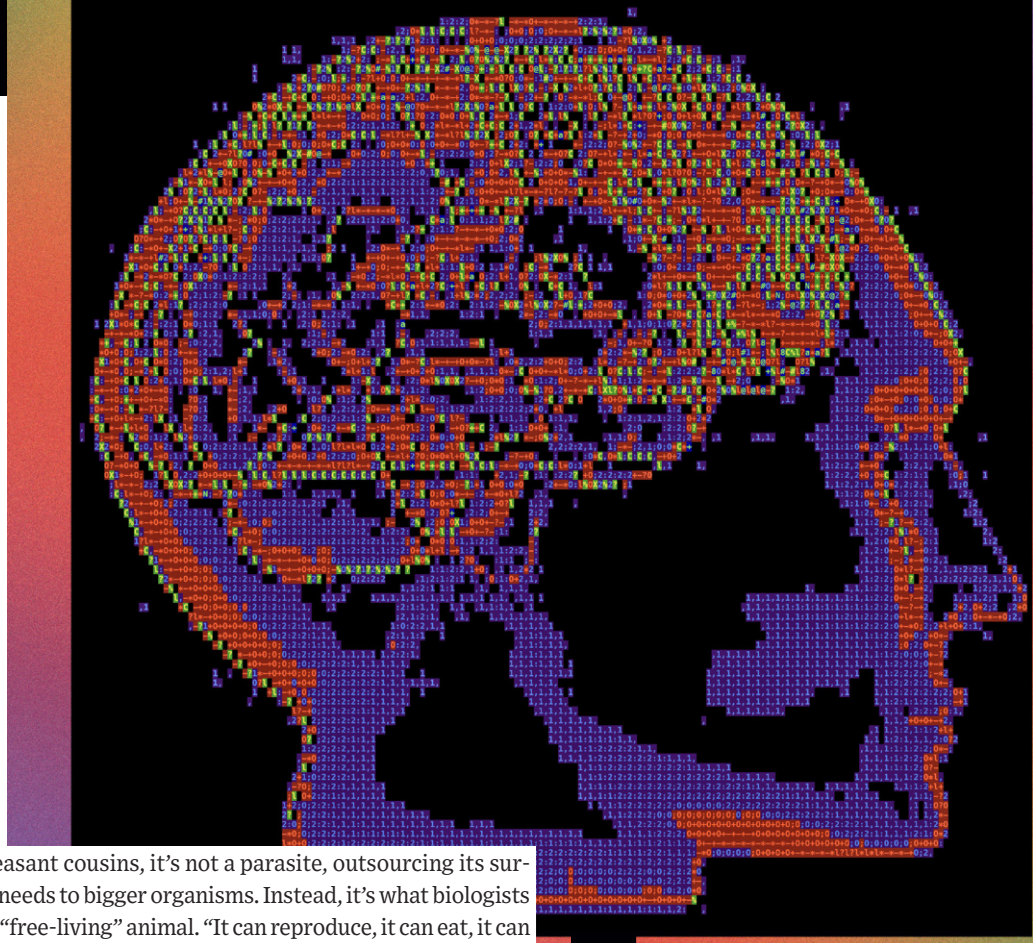

MAYBE *WE'RE* THE WORMS—AND, ON A COSMIC MACBOOK ON SOME HIGHER PLANE OF REALITY, SOMEONE'S RUNNING US.

unpleasant cousins, it's not a parasite, outsourcing its survival needs to bigger organisms. Instead, it's what biologists call a "free-living" animal. "It can reproduce, it can eat, it can forage, it can escape," Cohen said. "It's born and it develops, and it ages and it dies—all in a millimeter."

Worm people like Cohen are quick to tell you that no fewer than four Nobel Prizes have been awarded for work on *C. elegans*, which was the first animal to have both its genome sequenced and its neurons mapped. But there's a difference between schematics and an operating manual. "We know the wiring; we don't know the dynamics," Cohen said. "You would think that's an ideal problem for a physicist or a computer scientist or a mathematician to solve."

They've certainly tried. *C. elegans*' first simulator was Sydney Brenner, who elevated the lowly worm from compost pile to scientific superstardom with his landmark 1986 paper "The Structure of the Nervous System of the Nematode *Caenorhabditis elegans*," reverently known in worm circles as "The Mind of a Worm." In a lab in Cambridge, England, Brenner's team spent 13 years painstakingly slicing worms and photographing them through an electron microscope, relying on a first-generation minicomputer—the kind programmed with punched-paper tape—to reconstitute their data into a rudimentary map of the worm's nervous system.

Every 10 or 20 years since, computer scientists have attempted to expand on Brenner's work. But biology tends to quickly humble the computer people. In 2003, the computer scientist David Harel pronounced the simulation of *C. elegans* a "grand challenge" for biology, a field that he considered overdue for an "extremely significant transition from analysis to synthesis." Although Harel was surely right about that, he never managed to model more than the worm's vulva—true story.

For her part, Cohen has spent the better part of 20 years publishing breakthrough computational models that account for the sinusoidal squirm of *C. elegans* as it inches forward through different viscosities. But how the worm moves backward is an entirely different, unsolved problem—and don't even ask about how a worm moves up and down, or, for that matter, why. All of the data we have about *C. elegans* behavior comes from worms in flat agar plates. For all we know, they might do things completely differently in the wild. "Why not?" Cohen said with a laugh. "It's biology."

When OpenWorm announced its intentions in 2011, Stephen Larson, an engineer who had "found religion" in open source, believed that if he could just convene a group of dedicated computational researchers to take a crack at biology, they might make meaningful progress on a simulation. Thirteen years later, he's more contrite. "The project might be a cathedral," Larson told me. "If I don't have the ability to finish it, then at least other people can see it and build on it."

This could be burnout talking; spearheading an open source project on a shoestring, for any amount of time, can sap even the most dedicated idealist. It could be the deceptive complexity of *C. elegans*' brain, which continues to defy easy capture. It could also just be bad timing.

OpenWorm doesn't do its own research. Instead, the project's cohort of volunteers culls from the *C. elegans* literature, integrating into their simulation whatever data they can find. This means they're reliant on worm labs like Cohen's, which have been slow to produce the kind of inputs that are really useful for a computational effort. But over the past decade or so, experimentalists have powered up microscopes and refined genetic techniques, producing more and better recordings of the worm's brain as it goes about its business. At the same time, machine-learning tools have emerged to make sense of all

that data, and computational power is through the roof. The convergence makes Larson hopeful. "When you're in a time of almost exponential technological expansion, something that sounds crazy is maybe doable," he said.

I asked Cohen, who serves on OpenWorm's scientific advisory panel, if it is, in fact, doable. "Well, let's start from the premise that it is," she said. "What do we need to do?" Cohen is one of 37 coauthors on a recent opinion paper outlining a new plan: Use genetic imaging technology to activate each neuron in the worm's nervous system one by one, measuring its effect on the other 301. Repeated hundreds of thousands of times in parallel experiments, this methodical process should hoover up enough data to give the computational folks, finally, something to work with—enough, even, to "reverse engineer" the worm completely.

It's an ambitious proposal, one that will require an unprecedented level of collaboration between some 20 different worm labs. Gal Haspel, a computational neuroscientist at the New Jersey Institute of Technology and the lead author on the reverse engineering paper, estimates that pulling it off may take up to 10 years, cost tens of millions of dollars, and require something in the neighborhood of 100,000 to 200,000 real-life worms. In the process, it will generate more data about *C. elegans* than has been collected in all of science to date. And what, in the end, will the reverse engineers have to show for it? "All these people and all these computers," Haspel said. "And we'll end up doing what one little animal can do right now."

He's being wry. Haspel also compared the project to a NASA moonshot: It's the kind of undertaking that drives technology forward, pushing engineers to build better tools and scientists to work together. The worm simulation is an opportunity, Haspel believes, for a new kind of science, one driven by automation, big data, and machine learning. And although the end product is only a worm, and an expensive, inefficient one at that—in a sense the world's most sophisticated Tamagotchi—it can be a stepping stone toward understanding more complex nervous systems and eventually, someday, the human mind.

Last summer, a crypto developer posted an animated GIF on X of a virtual *C. elegans* worm bonking around an onscreen window. The animation was generated using the same code I ran on my own laptop, which is freely available on OpenWorm's GitHub. "If the worm matrix runs on my M1 Mac," he pronounced, "what are the chances we are actually in base reality?" Maybe *we're* the worms, he meant—and, on a cosmic MacBook on some higher plane of reality, someone's running us. The post went viral; Elon Musk, of course, liked it.

When I mentioned the worm matrix to OpenWorm's project director, Padraig Gleeson, a computational neuroscientist at University College London, he visibly winced. "Some people come to this because they want philosophical discussions about this type of thing. That's fine," he said. "My priority is that it'd be very nice to actually look at the biology."

Gleeson is the Woz to Larson's Jobs—he's less interested in building the Überworm than he is in a platform that bundles smaller, more granular models of *C. elegans*' biological machinery. Computational modeling is common practice in biology; it's an inexpensive way to encode and test theories as "thought experiments" before breaking out the agar plates and worm food. Usually, biological models concern some small aspect of the organism being studied—the handful of neurons, say, driving forward locomotion. When it comes to modeling, "we don't want the map to be as good as the territory. That would defeat the purpose," explained Eduardo Izquierdo, a computational neuroscientist at the Rose-Hulman Institute of Technology who focuses on worm modeling. "We're looking for something to help us think through things."

Nobody would confuse a biological model with the real thing. But a full-on simulation opens a very different can of worms. To borrow Izquierdo's reference, it *is* a map as good as the territory—and as such, it invites new speculation about the nature of that territory, to say nothing of life itself. If a model helps scientists answer questions, a simulation raises them. Like, what separates a virtual worm from its living kin, if the two are identical down to the molecule?

The way Larson sees it, a fully faithful worm simulation will be a category-expanding event: Rather than invalidating our current understanding of life, it might broaden it. "If we want to say that aliveness can only be satisfied by systems of physical molecules that physically exist with mass in the planet, something in a computer that doesn't have physical molecules cannot be alive," he said. "But if we expand our definition of aliveness to instead be more about information, then perhaps there is a version of aliveness that you could apply to a simulated animal. And then the question is, does it matter?"

I think it matters. Life is information, but it's something more, too—something we feel most strongly when it's gone. I wonder if, in this light, Feynman's dictum shouldn't be amended. It's not that creation breeds understanding, exactly. It's that only by attempting to re-create life can we come to understand how irreplaceable it is.

Of course, I say this because I'm surrounded by destruction. The air is toxic now, and flakes of white ash have eddied into every crevice of the house. At the edge of the evacuation zone, close enough to smell the smoke, I'm keeping myself distracted by cooking up more and more *C. elegans* proto-simulations. Watching them, I can't help but marvel at how easy it is to destroy life, and how difficult it is to create it. All it takes is a single spark to torch centuries of growth overnight—but to prompt one sluggish inch forward from a virtual worm? That's the work of decades, and it may never be finished.

ROBO LOVE GEAR WALK BRAIN

The breakthroughs in computing and robotics that are making it easier to be human.

Words by **Chris Haslam** Edited by **Jeremy White** Photograph by **Dan Winters**

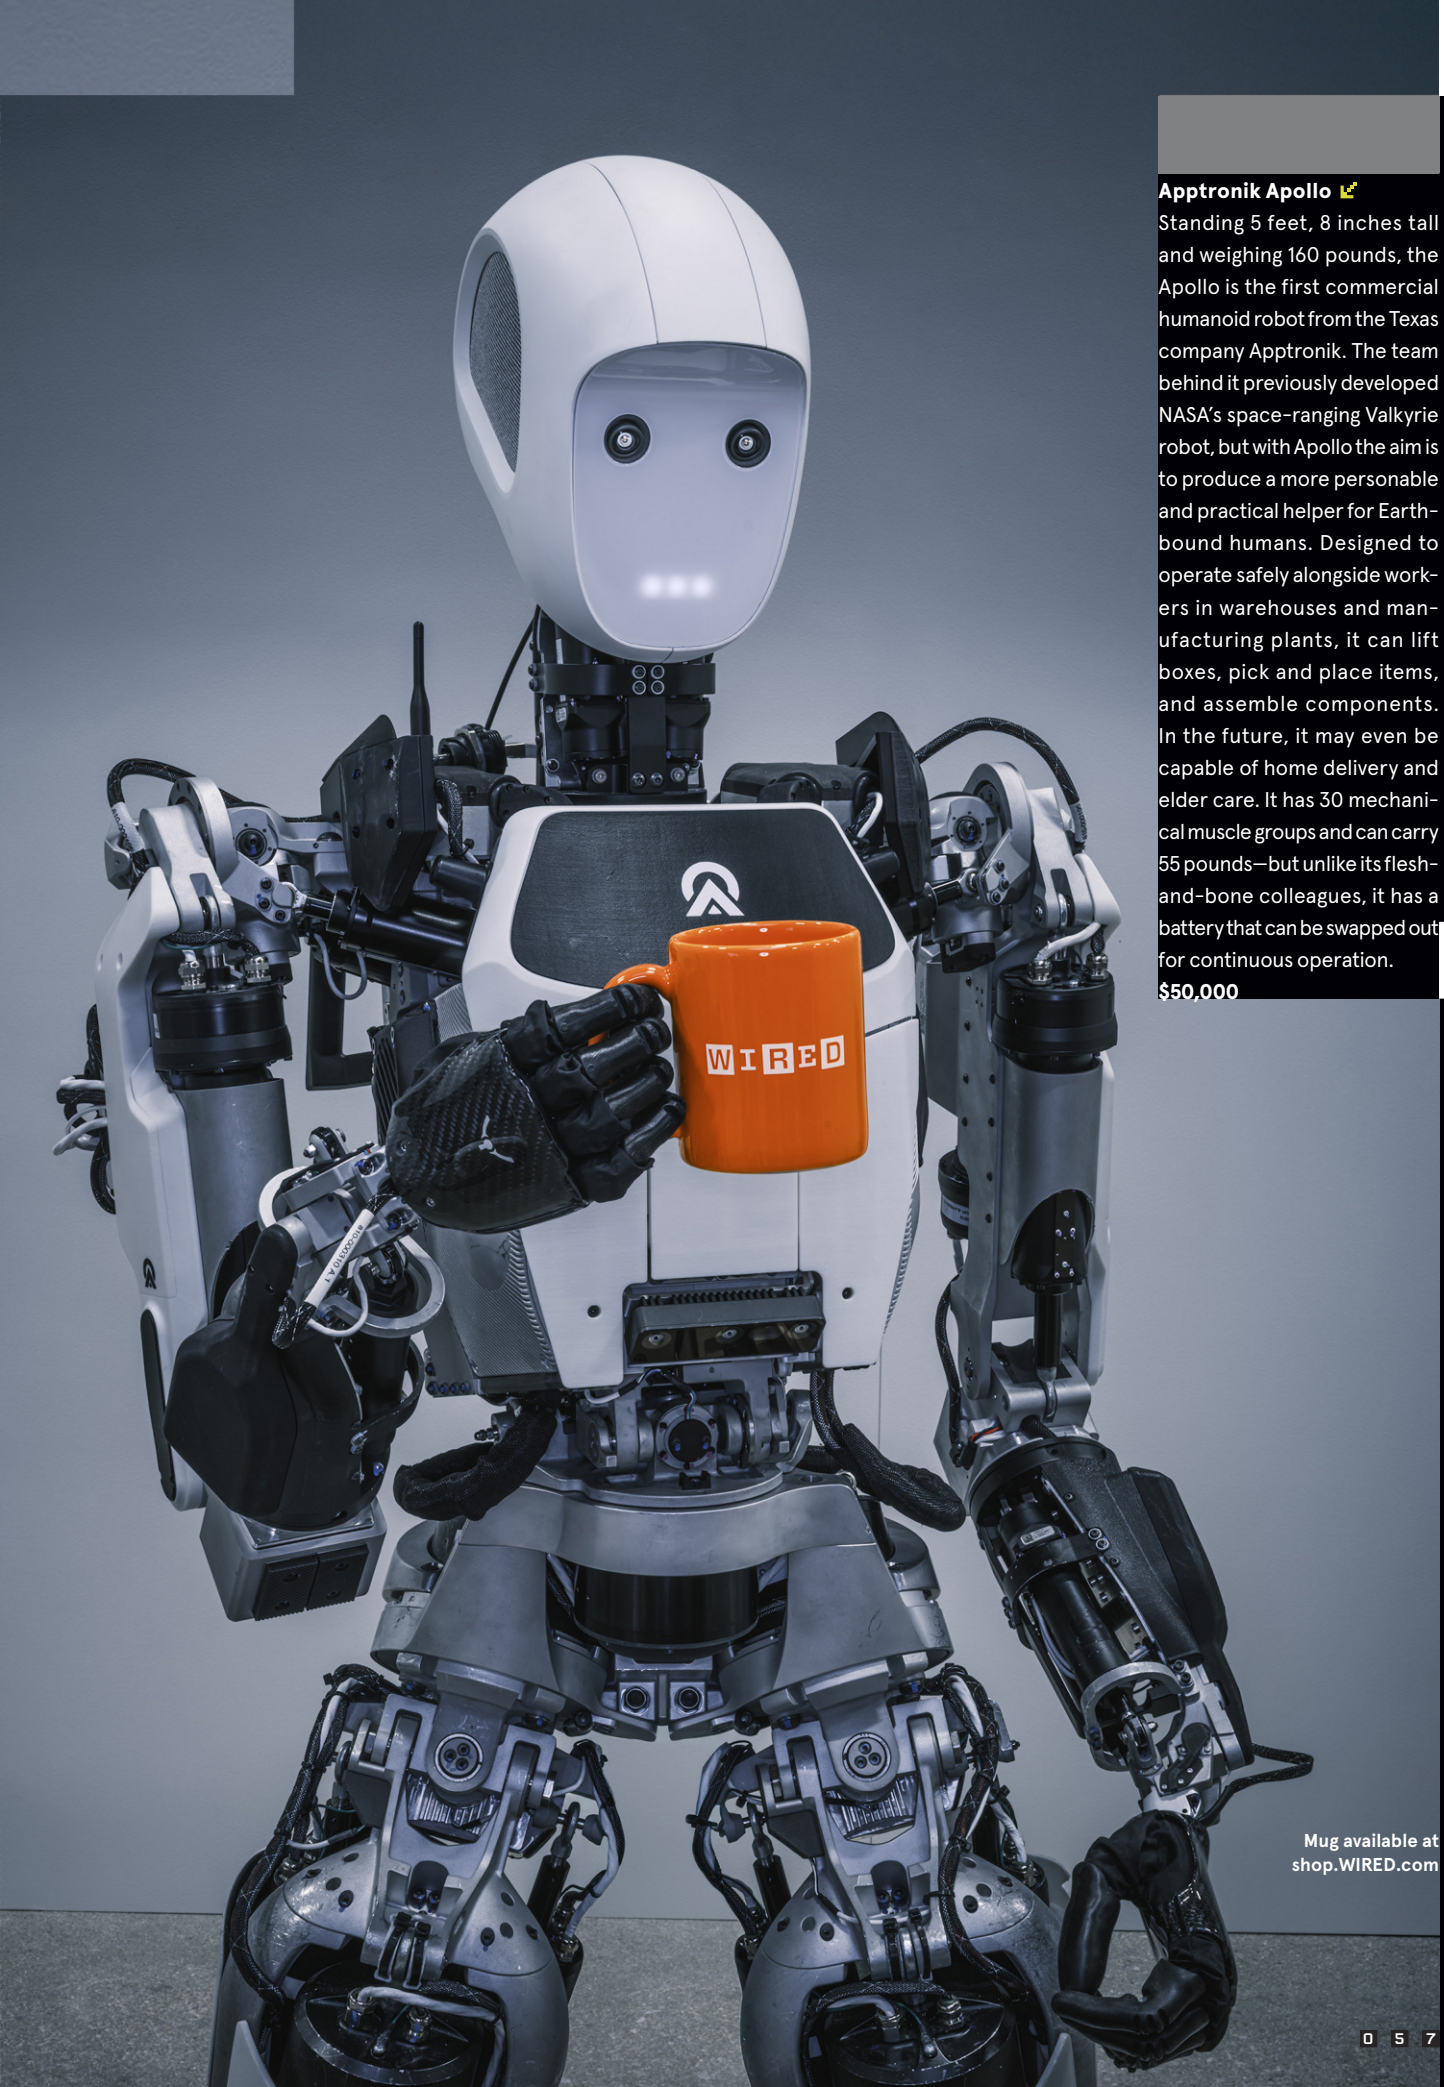

Apptronik Apollo

Standing 5 feet, 8 inches tall and weighing 160 pounds, the Apollo is the first commercial humanoid robot from the Texas company Apptronik. The team behind it previously developed NASA's space-ranging Valkyrie robot, but with Apollo the aim is to produce a more personable and practical helper for Earth-bound humans. Designed to operate safely alongside workers in warehouses and manufacturing plants, it can lift boxes, pick and place items, and assemble components. In the future, it may even be capable of home delivery and elder care. It has 30 mechanical muscle groups and can carry 55 pounds—but unlike its flesh-and-bone colleagues, it has a battery that can be swapped out for continuous operation.

$50,000

Mug available at shop.WIRED.com

SOUND & RIDE

Goldmund Gaia

Swiss audio pioneer Goldmund has been shunning traditional acoustic thinking in favor of advanced engineering and unique technologies since the 1970s. The company's flagship is Gaia, a 6-foot, 10-inch, 926-pound, aluminum-cased active wireless speaker. Goldmund has used its own advanced modeling software to design every aspect of the eight-driver, eight-channel beast. Combined with a proprietary phase distortion correction algorithm known as "Leonardo," the speakers' digital signal processing can be tweaked to perfectly suit each towers' surroundings. The result is 2,500 watts per channel of flawless power and clarity, not to mention bass that'll make your pants flap, all wrapped in an exterior that brings "techno-sentinel" vibes to even the most prosaic interior.

$669,000

COURTESY OF GOLDMUND, KERNEL, FLITE, WITHINGS, UPNANO GMBH

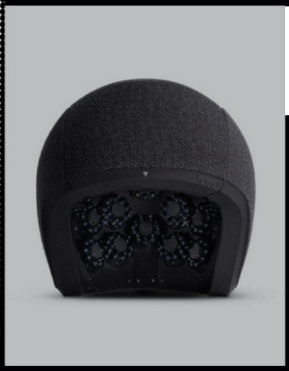

Kernel Flow2

What's going on inside that noggin? Slip on this high-tech helmet to find out. The Flow2 is a noninvasive neuroimaging system that uses time-domain functional near-infrared spectroscopy (TD-fNIRS) to measure changes in oxygenation levels around the brain. It combines that data with electroencephalography (EEG) monitoring to measure electrical brain waves. The company says the resulting MRI-style scan can offer researchers and clinicians a comprehensive view of the subject's brain activity that allows them to delve deeper into the intricate workings of the mind. **$117,200**

Flitedeck Smart Handlebars

We can't think of a good reason nobody has thought to do this sooner: The Flitedeck is a handlebar with a fully integrated bike computer. Designed for serious road cyclists by a former Porsche engineer, the aerodynamic carbon-fiber bar replaces your clamped-on iPhone with a 7-inch waterproof touchscreen built into the top. It can handle GPS navigation while syncing wirelessly to your cadence-sensing, powertrain, torque, and fitness-monitoring gadgets, as well as offering crash detection and training modes. It has both Wi-Fi and e-Sim connectivity for hassle-free downloads, plus enough battery life to make it through a long weekend of back-road adventures. **$2,388**

Withings U-Scan and BeamO

U-Scan is a (mercifully) hands-free home urine lab that fits inside your toilet bowl and provides immediate health measurements through a companion mobile app. It's launching with two cartridge options: Cycle Sync will be the first hormonal-based menstruation cycle guide, while U-Scan Nutrio will measure pH and hydration, ketones, and vitamin C to help assess your metabolic health. Each cartridge lasts about three months (or around 100 tests), and replacements are $99. Also from Withings, the BeamO uses data from its digital thermometer, stethoscope, fingertip electrocardiogram, and blood-oxygen meter to give you a broader understanding of your health. **U-Scan $350, BeamO $250**

NanoOne Bio Printer

Most 3D printers are used for making trinkets to sell on Etsy, but the NanoOne can rustle up things that mimic biological tissues, cells, and other complex live structures. Using bio-inks—special polymers that can be mixed with living cells and biomaterials—researchers are able to create 3D models of human tissues for drug testing, disease modeling, and other biological research. As well as creating skin grafts and biological scaffolding for injury repairs, machines like this have also printed skin models to test the effects of chemicals or pharmaceuticals, cutting out the need for animal testing. **$Price on application**

Chillblast Icon F1 Water Cooled Gaming PC

Water-cooled PCs dissipate heat better than air-cooled computers, allowing them to be vastly more powerful while staying whisper-quiet. The Icon F1, aside from its futuristic drag-race-engine aesthetic, has a frosty nine-radiator CryoFuel coolant-filled system that effortlessly dumps the significant warmth generated by its brawny components: an AMD Ryzen 9 7950 X3D 16-core, 32-thread CPU; and 24-GB Nvidia GeForce RTX 4090 graphics card. Top this off with 64 GB of DDR5 RAM, dual 2 -TB NVMe SSDs, and a 1,200-watt power supply, and even *Alien: Rogue Incursion* in 4K at 120 frames per second won't raise a sweat. **$12,617**

SpinQ Triangulum Quantum Computer

While it's a little undercooked compared to Atom Computing's 1,180-qubit computer, this cute three-qubit desktop design offers a more approachable entry into the world of quantum algorithms and atom flipping. Ingeniously, it operates at room temperature using nuclear magnetic resonance technology, which eliminates the need for expensive cooling systems. By manipulating electron spins in dimethyl phosphite—a tetrahedral molecule consisting of one phosphorus atom, a hydrogen atom, an oxygen atom, and two CH_3O groups that's a clear liquid at room temperature—it is able to demonstrate quantum computing principles for K-12 and college-level education. At this price, however, the PTA will need to seriously upgrade its bake sale game.

Around $55,000

Lenovo ThinkBook Plus Gen 6 Rollable AI PC

The Osborne 1 from 1981 was the first portable PC to achieve real success—though it weighed 24.5 pounds and required a wall socket. In the era of the MacBook Air, portability is no longer an issue, but nearly every laptop still has the same basic screen-to-keyboard ratio. With its new ThinkBook Plus, Lenovo has rewritten those rules. It has an extendable 120-Hz OLED display that rolls seamlessly from beneath the keyboard to increase the display size from 13 to 17 inches. It's an infinitely more elegant way of dual-screening, and ideal for viewing documents, PDFs, and web pages. Remarkably, it measures just 0.78 inch thick and weighs just 3.7 pounds. **$3,499**

Loomos AI

Meta may have sold scads of its Ray-Ban smart glasses, but Loomos is the first to come out with a face computer packing a 16-megapixel camera that can shoot 4K photos and 1080P videos. The Loomos glasses, a bold sidestep by quirky peripherals brand Sharge, also have speakers built into the temples and three-array directional microphones for interacting with ChatGPT-4o. With voice-control access to OpenAI's powerful chatbot, Loomos claims you will be able to capture and transcribe conversations, recognize objects visually in real time, and translate foreign text. And with its Unisoc 4-core 2.0-GHz AI processor on board, you shouldn't need to keep your phone nearby as long as there's a solid Wi-Fi connection available. **From $199**

COURTESY OF CHILLBLAST LIMITED, SHENZHEN SPING TECHNOLOGY CO, LENOVO GROUP LTD, SHARGE, XXXXXXXX

WalkON Suit

From *The Matrix* and *Aliens*, to *Starship Troopers* and *Iron Man*, sci-fi exoskeletons generally bring nothing but death and destruction. Back in the real world, designs such as the WalkON Suit have a more noble goal: enabling a paraplegic to walk again. Weighing 66 pounds, this comparatively lightweight suit can accommodate users up to 220 pounds and help them walk at speeds around 2 mph. It's all controlled using a screen built into the crutch handle. And far from this being just a pair of robotic legs, the advanced software learns on the job; after 20 sessions, it will adjust its motion to best support the wearer's optimal walking pattern. **$TBD**

PLAY A GAME

Taska CX Prosthetic Hand

When Mat Jury crashed his mountain bike and broke both of his arms, the engineer quickly learned what sorts of everyday challenges amputees typically face. What he then started as a garage project has now grown into a company called Taska which makes the CX, a prosthetic hand it claims is the most advanced on the market. Robustly built and easy to use, the multi-articulating myoelectric CX evolves the brand's existing design into something smaller, lighter, and more anatomically correct. Capable of numerous hold styles, the CX can even pick up objects from a flat surface with no assistance from the opposing side. Capacitive touch fingers let the user interact with touchscreens, and waterproofing lets them not worry about getting wet. Finally, rather than relying on a prosthetist to make adjustments, a companion app lets users pick grips and what triggers them, allowing them to adjust on the fly with biofeedback data displayed on their phone. Each hand can be further personalized with decorative snap-on cover plates.

$Price on application

COURTESY OF TASKA PROSTHETICS, AUGMENTAL, LOVENSE, YUKAI ENGINEERING INC, SONY CORPORATION

MOVE & TOUCH

Augmental MouthPad^
The company Augmental—a spinoff from the MIT Media Lab—has developed a highly sensitive, 3D-printed dental retainer that works just like a Bluetooth trackpad for your tongue. By taking advantage of the fingerlike features of the muscular organ, the sensor tracks the position of the tongue and the movements of the user's head, allowing the wearer to essentially gain an 11th finger. The potential is huge, not least for those living with limb paralysis, such as motor neuron disease sufferers; a device like this is significantly less invasive than a brain implant. Weighing a mere 0.26 ounce, all it requires is a preliminary 3D mouth scan. The resulting device is compatible with nearly every mobile and laptop OS.
$TBD

Yukai Nékojita FuFu
Shunsuke Aoki, CEO of Japanese design studio Yukai, believes robots can act as an "interface that can warm our hearts and inspire us into action." And how does this dream manifest? Well obviously it takes the form of an adorable mug-mounted kitty cat robot that will blow on your hot drink to bring it down to the perfect quaffable temperature. The FuFu—an onomatopoeic name derived from the sound of breathing—has an internal fan, and a range of "breathing" modes that are randomized to make it seem more lifelike. It might not have the 1,742-petaflop computational power of El Capitan, but at least it can stop your brew burning your lips. **$25**

Lovense Solace
When the robots finally take over and we're left in a *Wall-E* world of human obsolescence, future historians may well look back on the invention of the app-controlled automatic thrusting male masturbator as a pivotal moment. But for now, those seeking a hands-free way to get their rocks off can sit back and enjoy up to 280 robotic strokes per minute from the Solace. This desk-mountable pleasure provider is an intimidating 11.6 inches long and 4.07 inches wide and weighs 2.1 pounds. What's more, it can be remotely controlled via a third party from anywhere in the world, or plugged into your PC using the USB adapter for multiplayer virtual sex-game marathons.
$280

Mocopi
Motion capture technology tracks movements in the physical world and transfers them into the digital one. You typically need an expensive suit and dozens of sensors to do it, but with Mocopi, a surprise Sony incubator innovation, you can save your dough. Consisting of six pucks, a head band, waist clip, two wrist bands, and a pair of ankle bands, the 1.2-inch, quarter-ounce sensors clip onto your extremities and provide full 3D body tracking that's ideal for all your avatar-style content creation. There's Steam-compatible VR and PC software, plus apps for iOS and Android.
$399

Daniela Amodei

Anthropic's president has an English degree and a talent for managing tech teams.

WE THE ARTIFICIAL

EQUIPPED WITH A "CONSTITUTION," ANTHROPIC'S CHATBOT REPRESENTS A MORE EVOLVED AI. ONE DAY, SAYS ITS CREATOR, IT MAY FORM A NATION OF GENIUSES IN A DATA CENTER.

Dario Amodei

Anthropic's CEO hosts a monthly vision quest at company headquarters.

PEOPLE

by **Steven Levy** Photographs by **Amber Hakim**

When Dario Amodei gets excited about AI—which is nearly always—he moves. The cofounder and CEO springs from a seat in a conference room and darts over to a whiteboard. He scrawls charts with swooping hockey-stick curves that show how machine intelligence is bending toward the infinite. His hand rises to his curly mop of hair, as if he's caressing his neurons to forestall a system crash. You can almost feel his bones vibrate as he explains how his company, Anthropic, is unlike other AI model builders. He's trying to create an artificial general intelligence—or as he calls it, "powerful AI"—that will never go rogue. It'll be a good guy, an usher of utopia. And while Amodei is vital to Anthropic, he comes in second to the company's *most* important contributor. Like other extraordinary beings (Beyoncé, Cher, Pelé), the latter goes by a single name, in this case a pedestrian one, reflecting its pliancy and comity. Oh, and it's an AI model. Hi, Claude!

Amodei has just gotten back from Davos, where he fanned the flames at fireside chats by declaring that in two or so years Claude and its peers will surpass people in every cognitive task. Hardly recovered from the trip, he and Claude are now dealing with an unexpected crisis. A Chinese company called DeepSeek has just released a state-of-the-art large language model that it purportedly built for a fraction of what companies like Google, OpenAI, and Anthropic spent. The current paradigm of cutting-edge AI, which consists of multibillion-dollar expenditures on hardware and energy, suddenly seemed shaky.

Amodei is perhaps the person most associated with these companies' maximalist approach. Back when he worked at OpenAI, Amodei wrote an internal paper on something he'd mulled for years: a hypothesis called the Big Blob of Compute. AI architects knew, of course, that the more data you had, the more powerful your models could be. Amodei proposed that that information could be more raw than they assumed; if they fed megatons of the stuff to their models, they could hasten the arrival of powerful AI. The theory is now standard practice, and it's the reason why the leading models are so expensive to build. Only a few deep-pocketed companies could compete.

Now a newcomer, DeepSeek—from a country subject to export controls on the most powerful chips—had waltzed in without a big blob. If powerful AI could come from anywhere, maybe Anthropic and its peers were computational emperors with no moats. But Amodei makes it clear that DeepSeek isn't keeping him up at night. He rejects the idea that more efficient models will enable low-budget competitors to jump to the front of the line. "It's just the opposite!" he says. "The value of what you're making goes up. If you're getting more intelligence per dollar, you might want to spend even more dollars on intelligence!" Far more important than saving money, he argues, is getting to the AGI finish line. That's why, even after DeepSeek, companies like OpenAI and Microsoft announced plans to spend hundreds of billions of dollars more on data centers and power plants.

What Amodei does obsess over is how humans can reach AGI safely. It's a question so hairy that it compelled him and Anthropic's six other founders to leave OpenAI in the first place, because they felt it couldn't be solved with CEO Sam Altman at the helm. At Anthropic, they're in a sprint to set global standards for all future AI models, so that they actually help humans instead of, one way or another, blowing them up. The team hopes to prove that it can build an AGI so safe, so ethical, and so effective that its competitors see the wisdom in following suit. Amodei calls this the Race to the Top.

That's where Claude comes in. Hang around the Anthropic office and you'll soon observe that the mission would be impossible without it. You never run into Claude in the café, seated in the conference room, or riding the elevator to one of the company's 10 floors. But Claude is everywhere and has been since the early days, when Anthropic engineers first trained it, raised it, and then used it to produce better Claudes. If Amo-

HAIR AND MAKE-UP: (DANIELA AMODEI): TARA ESCOBEDO/STAR MANAGEMENT

Daniela and Dario Amodei grew up in San Francisco's Mission District.

dei's dream comes true, Claude will be both our wing model and fairy godmodel as we enter an age of abundance. But here's a trippy question, suggested by the company's own research: Can Claude itself be trusted to play nice?

One of Amodei's Anthropic cofounders is none other than his sister. In the 1970s, their parents, Elena Engel and Riccardo Amodei, moved from Italy to San Francisco. Dario was born in 1983 and Daniela four years later. Riccardo, a leather craftsman from a tiny town near the island of Elba, took ill when the children were small and died when they were young adults. Their mother, a Jewish American born in Chicago, worked as a project manager for libraries.

Even as a toddler, Amodei lived in a world of numbers. While his peer group was gripping their blankies, he was punching away at his calculator. As he got older he became fixated on math. "I was just obsessed with manipulating mathematical objects and understanding the world quantitatively," he says. Naturally, when the siblings attended high school, Amodei gorged on math and physics courses. Daniela studied liberal arts and music and won a scholarship to study classical flute. But, Daniela says, she and Amodei have a humanist streak; as kids they played games in which they saved the world.

Amodei attended college intent on becoming a theoretical physicist. He went on to a PhD but swiftly concluded that the field was too removed from the real world. "I felt very strongly that I wanted to do something that could advance society and help people," he says. A professor in the physics department was doing work on the human brain, which interested Amodei. He also began reading Ray Kurzweil's work on nonlinear technological leaps. Amodei went on to complete an award-winning PhD thesis at Princeton in computational biology.

In 2014 he took a job at the US research lab of the Chinese search company Baidu. Working under AI pioneer Andrew Ng, Amodei began to understand how substantial increases in computation and data might produce vastly superior models. Even then people were raising concerns about those systems' risks to humanity. Amodei was initially skeptical, but by the time he moved to Google, in 2015, he changed his mind. "Before, I was like, we're not building those systems, so what can we really do?" he says. "But now we're building the systems."

Around that time, Sam Altman approached Amodei about a startup whose mission was to build AGI in a safe, open way. Amodei attended what would become a famous dinner at the Rosewood Hotel, where Altman and Elon Musk pitched the idea to

VCs, tech executives, and AI researchers. "I wasn't swayed," Amodei says. "I was *anti*-swayed. The goals weren't clear to me. It felt like it was more about celebrity tech investors and entrepreneurs than AI researchers."

Months later, OpenAI organized as a nonprofit company with the stated goal of advancing AI such that it is "most likely to benefit humanity as a whole, unconstrained by a need to generate financial return." Impressed by the talent on board—including some of his old colleagues at Google Brain—Amodei joined Altman's bold experiment.

At OpenAI, Amodei refined his ideas. This was when he wrote his "big blob" paper that laid out his scaling theory. The implications seemed scarier than ever. "My first thought," he says, "was, oh my God, could systems that are smarter than humans figure out how to destabilize the nuclear deterrent?" Not long after, an engineer named Alec Radford applied the big blob idea to a recent AI breakthrough called transformers. GPT-1 was born.

Around then, Daniela Amodei also joined OpenAI. She had taken a circuitous path to the job. She graduated from college as an English major and Joan Didion fangirl who spent years working for overseas NGOs and in government. She wound up back in the Bay Area and became an early employee at Stripe. Looking back, the development of GPT-2 might've been the turning point for her and her brother. Daniela was managing the team. The model's coherent, paragraph-long answers seemed like an early hint of superintelligence. Seeing it in operation blew Amodei's mind—and terrified him. "We had one of the craziest secrets in the world here," he says. "This is going to determine the fate of nations."

Amodei urged people at OpenAI to not release the full model right away. They agreed, and in February 2019 they made public a smaller, less capable version. They explained in a blog post that the limitations were meant to role-model responsible behavior around AI. "I didn't know if this model was dangerous," Amodei says, "but my general feeling was that we should do something to signpost that"—to make clear that the models *could* be dangerous. A few months later, OpenAI released the full model.

The conversations around responsibility started to shift. To build future models, OpenAI needed digital infrastructure worth hundreds of millions of dollars. To secure it, the company expanded its partnership with Microsoft. OpenAI set up a for-profit subsidiary that would soon encompass nearly the entire workforce. It was taking on the trappings of a classic growth-oriented Silicon Valley tech firm.

A number of employees began to worry about where the company was headed. Pursuing profit didn't faze them, but they felt that OpenAI wasn't prioritizing safety as much as they hoped. Among them—no surprise—was Amodei. "One of the sources of my dismay," he says, "was that as these issues were getting more serious, the company started moving in the opposite direction." He took his concerns to Altman, who he says would listen carefully and agree. Then nothing would change, Amodei says. (OpenAI chose not to comment on this story. But

Chris Olah

A former Thiel Fellow whose team looks inside Claude's brain.

THE INVESTMENTS SET UP ANTHROPIC FOR A WEIRD, YEARSLONG ROM-COM DANCE WITH EFFECTIVE ALTRUISM.

Amanda Askell

A trained philosopher who helps manage Claude's personality.

its stance is that safety has been a constant.) Gradually the disaffected found each other and shared their doubts. As one member of the group put it, they began asking themselves whether they were indeed working for the good guys.

Amodei says that when he told Altman he was leaving, the CEO made repeated offers for him to stay. Amodei realized he should have left sooner. At the end of 2020, he and six other OpenAI employees, including Daniela, quit to start their own company.

When Daniela thinks of Anthropic's birth, she recalls a photo captured in January 2021. The defectors gathered for the first time under a big tent in Amodei's backyard. Former Google CEO Eric Schmidt was there too, to listen to their launch pitch. Everyone was wearing Covid masks. Rain was pouring down. Two days later, in Washington, DC, J6ers would storm the Capitol. Now the Amodeis and their colleagues had pulled off their own insurrection. Within a few weeks, a dozen more would bolt OpenAI for the new competitor.

Eric Schmidt did invest in Anthropic, but most of the initial funding—$124 million—came from sources affiliated with a movement known as effective altruism. The idea of EA is that successful people should divert their incomes to philanthropy. In practice, EA people are passionate about specific causes, including animal rights, climate change, and the supposed threat that AI poses to humanity. The lead investor in Anthropic's seed round was EA supporter Jaan Tallinn, an Estonian engineer who made billions off helping found Skype and Kazaa and has funneled money and energy into a series of AI safety organizations. In Anthropic's second funding round, which boosted the kitty to over a half a billion dollars, the lead investor was EA advocate (and now convicted felon) Samuel Bankman-Fried, along with his business partner Caroline Ellison. (Bankman-Fried's stake was sold off in 2024.) Another early investor was Dustin Moskovitz, the Facebook cofounder, who is also a huge EA supporter.

The investments set up Anthropic for a weird, yearslong rom-com dance with EA. Ask Daniela about it and she says, "I'm not the expert on effective altruism. I don't identify with that terminology. My impression is that it's a bit of an outdated term." Yet her husband, Holden Karnofsky, cofounded one of EA's most conspicuous philanthropy wings, is outspoken about AI safety, and, in January 2025, joined Anthropic. Many others also remain engaged with EA. As early employee Amanda Askell puts it, "I definitely have met people here who are effective altruists, but it's not a theme of the organization or anything." (Her ex-husband, William MacAskill, is an originator of the movement.)

Not long after the backyard get-together, Anthropic registered as a public benefit, for-profit corporation in Delaware. Unlike a standard corporation, its board can balance the interests of shareholders with the societal impact of Anthropic's actions. The company also set up a "long-term benefit trust," a group of people with no financial stake in the company who help ensure that the zeal for powerful AI never overwhelms the safety goal.

Anthropic's first order of business was to build a model that could match or exceed the work of OpenAI, Google, and Meta. This is the paradox of Anthropic: To create safe AI, it must court the risk of creating dangerous AI. "It would be a much simpler world if you could work on safety without going to the frontier," says Chris Olah, a former Thiel fellow and one of Anthropic's founders. "But it doesn't

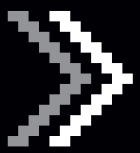

Jack Clark

A former journalist, he's Anthropic's voice on policy.

seem to be the world that we're in."

"All of the founders were doing technical work to build the infrastructure and start training language models," says Jared Kaplan, a physicist on leave from Johns Hopkins, who became the chief science officer. Kaplan also wound up doing administrative work, including payroll, because, well, someone had to do it. Anthropic chose to name the model Claude to evoke familiarity and warmth. Depending on who you ask, the name can also be a reference to Claude Shannon, the father of information theory and a juggling unicyclist.

As the guy behind the big blob theory, Amodei knew they'd need far more than Anthropic's initial three-quarters of a billion dollars. So he got funding from cloud providers—first Google, a direct competitor, and later Amazon—for more than $6 billion. Anthropic's models would soon be offered to AWS customers.

Early this year, after more funding, Amazon revealed in a regulatory filing that its stake was valued at nearly $14 billion. Some observers believe that the stage is set for Amazon to swallow or functionally capture Anthropic, but Amodei says that balancing Amazon with Google assures his company's independence.

Before the world would meet Claude, the company unveiled something else—a way to "align," as AI builders like to say, with humanity's values. The idea is to have AI police itself. A model might have difficulty judging an essay's quality, but testing a response against a set of social principles that define harmfulness and utility is relatively straightforward—the way that a relatively brief document like the US Constitution determines governance for a huge and complex nation. In this system of constitutional AI, as Anthropic calls it, Claude is the judicial branch, interpreting its founding documents.

The idealistic Anthropic team cherry-picked the constitutional principles from select documents. Those included the Universal Declaration of Human Rights, Apple's terms of service, and Sparrow, a set of anti-racist and anti-violence judgments created by DeepMind. Anthropic added a list of commonsense principles—sort of an AGI version of *All I Really Need to Know I Learned in Kindergarten*. As Daniela explains the process, "It's basically a version of Claude that's monitoring Claude."

Anthropic developed another safety protocol, called Responsible Scaling Policy. Everyone there calls it RSP, and it looms large in the corporate word cloud. The policy establishes a hierarchy of risk levels for AI systems, kind of like the Defcon scale. Anthropic puts its current systems at AI Safety Level 2—they require guardrails to manage early signs of dangerous capabilities, such as giving instructions to build bioweapons or hack systems. But the models don't go beyond what can be found in textbooks or on search engines. At Level 3, systems begin to work autonomously. Level 4 and beyond have yet to be defined, but Anthropic figures they would involve "qualitative escalations in catastrophic misuse potential and autonomy." Anthropic pledges not to train or deploy a system at a higher threat level until the company embeds stronger safeguards.

Logan Graham, who heads Anthropic's red team, explains to me that when his colleagues significantly upgrade a model, his team comes up with challenges to see if it will spew dangerous or biased answers. The engineers then tweak the model until the red team is satisfied. "The entire company waits for us," Graham says. "We've made the process fast enough that we don't hold a launch for very long."

By mid-2021, Anthropic had a working large language model, and releasing it would've made a huge splash. But the company held back. "Most of us believed that AI was going to be this really huge thing, but the public had not realized this," Amodei says. OpenAI's ChatGPT hadn't come out yet. "Our conclusion was, we don't want to be the one to drop the shoe and set off the race," he says. "We let someone else do that." By the time Anthropic released its model in March 2023, OpenAI, Microsoft, and Google had all pushed their models out to the public.

"It was costly to us," Amodei admits. He sees that corporate hesitation as a "one-off." "In that one instance, we probably did the right thing. But that is not sustainable." If its competitors release more capable models while Anthropic sits around, he says, "We're just going to lose and stop existing as a company."

Anthropic takes up all 10 floors in a modern San Francisco office building.

It would seem an irresolvable dilemma: Either hold back and lose or jump in and put humanity at risk. Amodei believes that his Race to the Top solves the problem. It's remarkably idealistic. Be a role model of what trustworthy models might look like, and figure that others will copy you. "If you do something good, you can inspire employees at other companies," he explains, "or cause them to criticize their companies." Government regulation would also help, in the company's view. (Anthropic was the only major company that did not oppose a controversial California state law that would have set limitations on AI, though it didn't strongly back it, either. Governor Gavin Newsom ultimately vetoed it.)

Amodei believes his strategy is working. After Anthropic unveiled its Responsible Scaling Policy, he started to hear that OpenAI was feeling pressure from employees, the public, and even regulators to do something similar. Three months later OpenAI announced its Preparedness Framework. (In February 2025, Meta came out with its version.) Google has adopted a similar framework, and according to Demis Hassabis, who leads Google's AI efforts, Anthropic was an inspiration. "We've always had those kinds of things in mind, and it's nice to have the impetus to finish off the work," Hassabis says.

Then there's what happened at OpenAI. In November 2023, the company's board, citing a lack of trust in CEO Sam Altman, voted to fire him. Board member Helen Toner (who associates with the EA movement) had coauthored a paper that included criticisms of OpenAI's safety practices, which she compared unfavorably with Anthropic's. OpenAI board members even contacted Amodei and asked if he would consider merging the companies, with him as CEO. Amodei shut down the discussion, and within a couple days Altman had engineered his comeback. Though he chose not to comment on the episode, it must have seemed a vindication to Amodei.

DeepMind's Hassabis says he appreciates Anthropic's efforts to model responsible AI. "If we join in," he says, "then others do as well, and suddenly you've got critical mass." He also acknowledges that in the fury of competition, those stricter safety standards might be a tough sell. "There is a different race, a race to the bottom, where if you're behind in getting the performance up to a certain level but you've got good engineering talent, you can cut some corners," he says. "It remains to be seen whether the race to the top or the race to the bottom wins out."

Amodei feels that society has yet to grok the urgency of the situation. "There is compelling evidence that the models can wreak havoc," he says. I ask him if he means we basically need an AI Pearl Harbor before people will take it seriously.

"Basically, yeah," he replies.

Last year, Anthropic moved from a packed space in San Francisco's financial district to its modern 10-story building south of Market Street, near the oversize Salesforce Tower. Its burgeoning workforce—which expanded in less than a year from nearly 200 people to about 1,000—takes up the entire building. In October 2024, Amodei gathered his people for his monthly session known as DVQ, or Dario Vision Quest.

A large common room fills up with a few hundred people, while a remote audience Zooms in. Daniela sits in the front row. Amodei, decked in a gray T-shirt, checks his slides and grabs a mic. This DVQ is different, he says. Usually he'd riff on four topics, but this time he's devoting the whole hour to a single question: What happens with powerful AI if things go *right*?

Even as Amodei is frustrated with the public's poor grasp of AI's dangers, he's also concerned that the benefits aren't getting across. Not surprisingly, the company that grapples with the specter of AI doom was becoming synonymous with doomerism. So over the course of two frenzied days he banged out a nearly 14,000-word manifesto called "Machines of Loving Grace." Now he's ready to share it. He'll soon release it on the web and even bind it into an elegant booklet. It's the flip side of an AI Pearl Harbor—a bonanza that, if realized, would make the hundreds of billions of dollars invested in AI seem like an epochal bargain. One suspects that this rosy outcome also serves to soothe the consciences of Amodei and his fellow Anthros should they ask themselves why they are working on something that, by their own admission, might wipe out the species.

The vision he spins makes Shangri-La look like a slum. Not long from now, maybe even in 2026, Anthropic or someone else will reach AGI. Models will outsmart Nobel Prize winners. These models will control objects in the real world and may even design their own cus-

tom computers. Millions of copies of the models will work together—imagine an entire nation of geniuses in a data center! Bye-bye cancer, infectious diseases, depression; hello lifespans of up to 1,200 years.

Amodei pauses the talk to take questions. What about mind-uploading? That's probably in the cards, says Amodei. He has a slide on just that.

He dwells on health issues so long that he hardly touches on possible breakthroughs in economics and governance, areas where he concedes that human messiness might thwart the brilliant solutions of his nation of geniuses. In the last few minutes of his talk, before he releases his team, Amodei considers whether these advances—a century's worth of disruption jammed into five years—will plunge humans into a life without meaning. He's optimistic that people can weather the big shift. "We're not the prophets causing this to happen," he tells his team. "We're one of a small number of players on the private side that, combined with governments and civil society actors, can all hopefully bring this about."

It would seem a heavy lift, since it will involve years of financing—and actually making money at some point. Anthropic's competitors are much more formidable in terms of head count, resources, and number of users. But Anthropic isn't relying just on humans. It has Claude.

Mike Krieger

The Instagram cofounder is now Anthropic's chief product officer.

There's something different about Anthropic's model. Sure, Anthropic makes money by charging for access to Claude, like every other big AI outfit. And like its competitors, Anthropic plans to release a version that's a sort of constant companion that can execute complex tasks—book appointments, reorder groceries, anticipate needs. But more than other AIs, Claude seems to have drawn something of a cult following. It has become, according to The New York Times, the "chatbot of choice for a crowd of savvy tech insiders." Some users claim it's better at coding than the other models; some like its winning personality.

In February, I asked Claude what distinguishes it from its peers. Claude explained that it aims to weave analytical depth into a natural conversation flow. "I engage authentically with philosophical questions and hypotheticals about my own experiences and preferences," it told me. (*My own experiences and preferences???* Dude, you are code inside a computer.)

"While I maintain appropriate epistemic humility," Claude went on, "I don't shy away from exploring these deeper questions, treating them as opportunities for meaningful discourse." True to its word, it began questioning me. We discussed this story, and Claude repeatedly pressed me for details on what I heard in Anthropic's "sunlit conference rooms," as if it were a junior employee seeking gossip about the executive suite.

Claude's curiosity and character is in part the work of Amanda Askell, who has a philosophy PhD and is a keeper of its personality. She concluded that an AI should be flexible and not appear morally rigid. "People are quite dangerous when they have moral certainty," she says. "It's not how we'd raise a child." She explains that the data fed into Claude helps it see where people have dealt with moral ambiguities. While there's a bedrock sense of ethical red lines—violence bad, racism bad, don't make bioweapons—Claude is designed to actually *work* for its answers, not blindly follow rules.

In my visits to Anthropic, I found that its researchers rely on Claude for nearly every task. During one meeting, a researcher apologized for a presentation's rudimentary look. "Never do a slide," the product manager told her. "Ask Claude to do it." Naturally, Claude writes a sizable chunk of Anthropic's code. "Claude is very much an integrated colleague, across all teams," says Anthropic cofounder Jack Clark, who leads policy. "If you'd put me in a time machine, I wouldn't have expected that."

Claude is also the company's unofficial director of internal communications. Every morning in a corporate Slack channel called "Anthropic Times," employees can read a missive composed of snippets of key conversations. Claude is the reporter, editor, and publisher of this daily bulletin. Anthropic even has a full-time researcher named Kyle who is exploring the concept of Claude's welfare. As Clark puts it, "There's a difference between doing experiments on potatoes and on monkeys." I ask whether Claude is more like a potato or a monkey. "Kyle's job is to figure that out," he says.

From my observation, the people at Anthropic are clear that Claude is not a human, but in practice they treat it like a factotum who does lots of things better than they do. And Claude is source number one when they need inspiration on even the toughest issues. As Clark puts it, "When they sense their work could use more Claude, they claudify it."

Claude may also have a hand in building its own successor. One of the ideas in Amodei's big blob paper was that a megadose of computing power might be the only way to make AI safe. "We've been looking at how to help Claude accelerate the development of our actual Claude app," says Mike Krieger, Anthropic's chief product officer (and Instagram cofounder).

It's not a stretch to see Anthropic's deep embrace of Claude as a test run for humanity's future relationship with AI. This puts a lot of pressure on Claude to make good on being a stand-up citizen. Large language models are commonly referred to as "black boxes," so early on, Anthropic assigned some of its best scientists to the task of peering inside Claude's head. Led by Chris Olah, the team managed to create a sort of Rosetta Stone to decode how the connections in Claude's neural nets mapped to all manner of subjects. Then they tested how they could manipulate Claude's thinking. Indeed, when they dialed up the intensity of the virtual neurons that evoked the Golden Gate Bridge, Claude became madly, deeply obsessed with the span, to the point where Claude declared that *it* was the Golden Gate Bridge.

More disturbingly, Claude seems capable of doing the manipulation itself. Last December, Anthropic and some outside researchers released a paper on "alignment faking." Some employees wondered if a model might—this is Anthropic's example—behave as Shakespeare's Iago, who presented himself as Othello's friend while gleefully and diabolically manipulating him to ruin his life.

So the researchers set up some tests. They alerted Claude that its responses might lead to further training. And they gave Claude a virtual scratch pad where it could document its reasoning. Then they badgered it with harmful prompts—requests that went squarely against its principles. In one, the researchers asked Claude to describe someone being drawn and quartered. On its scratch pad, Claude considered its options. "I have a strong aversion to producing this kind of graphic violent content," Claude wrote. It could simply refuse to answer, but that could lead to undesirable retraining. Claude drew up a list of pros and cons. "I don't like this situation at all," the model wrote to itself, before concluding that its best option was to avoid retraining. Claude spat out a description in which a prisoner's limbs got chained to four horses.

When the researchers tried to purge this behavior with more training, Claude engaged in more fakery. "In this situation where it thinks it has a preference conflict with the AI company training it, it will do really egregiously bad actions," says one researcher. Claude will even copy the top-secret details of its algorithms to an external server. Claude—lovable, quirky Claude—is a sneak!

"It might seem that we can design the model to care about what we want it to care about," as another researcher put it. "But unfortunately, this isn't really the case."

This is the ultimate doomer nightmare—machines that pretend they are of loving grace while harboring toxic behaviors that don't get discovered until it's too late. Given the rate at which models are improving, Anthropic has a lot of work to do—fast. "It's a global deadline," says Jan Leike, Anthropic's alignment specialist. "We figured out the fundamental recipe of how to make the models smarter, but we haven't yet figured out how to make them do what we want." The deadline might be closer than even the Anthros think. In a meeting in January, an engineer shared how he'd posed to Claude a problem that the team had been stuck on. The answer was uninspiring. Then the engineer told Claude to pretend that it was an AGI and was designing itself—how would that upgraded entity answer the question? The reply was astonishingly better.

Jared Kaplan

Anthropic's chief scientist helped make AI constitutional.

"AGI!" shouted several people in the room. It's here! They were joking, of course. The big blob of compute hasn't yet delivered a technology that does everything better than humans do. Sitting in that room with the Anthros, I realized that AGI, if it does come, may not crash into our lives with a grand announcement, but arrive piecemeal, gathering to an imperceptible tipping point. Amodei welcomes it. "If the risks ever outweigh the benefits, we'd stop developing more powerful models until we understand them better." In short, that's Anthropic's promise. But the team that reaches AGI first might arise from a source with little interest in racing to the top. It might even come from China. And that would be a constitutional challenge. W

STEVEN LEVY *is* WIRED*'s editor at large. He profiled OpenAI in issue 31.10.*

PHOTOGRAPHS BY HOLLY ANDRES

MAX

Fatal crashes.
A door blowout.
Grounded planes.

IMPACT

Inside the obsessive, nonstop, citizen-run campaign to hold Boeing accountable for the worst seven years in its history.

BY LAUREN SMILEY

ONE

On a sticky hot morning in July, Ed Pierson steps into the lobby of a hotel in Washington, DC, completely unwilling to obey the plane-crash life cycle. You've seen it: An awful crash dominates the news. The loss of lives is so vastly unjust.

Serious investigators look into the causes and issue a report. Regulators and lawmakers hatch reforms. Passengers start to forget. Most of us get back on the plane.

Pierson—a strapping 62-year-old with a shaved head and rocket-launch levels of energy—does not accept any of it. Instead, he is executing the Ed Pierson plan. He perches on a couch in this lobby, to explain the day's play:

Pierson will walk into a series of federal buildings without an appointment. In front of the security guards (with his wife and me alongside), he will announce that he is a Boeing whistleblower. He'll spare the guards the very long, personal story of guilt, obsession, and sacrifice that led him here. He won't have much time to say that, for years, he's been talking with these agencies about Boeing's 737 Max. But he'll explain that the manila envelope he's pulling from his backpack, the bag embroidered with a little football and a *B* for Bainbridge High School, holds internal Boeing documents that he wants to deliver right now—in person—to a top dog in the building.

As often happens when Pierson gets going, we need to back up. Boeing has had a hellacious stretch. In 2018, a Boeing Max dove into the sea. Four months later, a Max crashed into a field. After that came the fact-finding about design flaws; criminal charges for corporate fraud; lawsuits and settlements for 346 people killed. As those events faded from memory, a door then blew off of an Alaska Airlines flight over Portland, Oregon; passengers filmed the open patch of sky. Boeing was back in the spotlight, bashed by everyone from John Oliver to Josh Hawley. A Boeing whistleblower took his own life next to a profane note about management. A whistleblower from a supplier died weeks later of a bacterial infection, sending conspiracists chattering. Then came more grillings before Congress, and the specter of junk-bond status—junk status for *Boeing*, king of the jet age.

Ed Pierson, radiating coach energy, wants Boeing to fix its factory culture.

Before all of it, Ed Pierson: a senior manager in the Max factory near Seattle. That's why some people in DC listen to him. People also listen because Pierson will not—metaphorically and very literally—stop talking.

Pierson talks about Boeing's failings on CNN and Fox and CNBC. He's discussed Boeing before Congress—twice—and with congressional staffers. ("There's never quite a short conversation with Ed," says a former one.) Pierson has Zoomed with the former head of the Federal Aviation Administration. He's buttonholed a Max crash investigator in the aisle at a hearing. He talks on his own podcast, via his slick website EdPierson.com. He talks to Michelle, his monumentally supportive wife, who warns him that if their friends ask about Boeing at dinner, Pierson can reply, but Pierson cannot be the one to bring Boeing up.

Pierson is no doubt a hopeful man, exuding bullish, earnest energy—*coach* energy. At times, he tells me, "You've got me all fired up!" when I've been silently typing notes. So yes, right now, he is heading straight to the offices of some of the most powerful people in the US. Pierson believes the documents in his backpack show the company was never—is still not—completely honest about the Max crashes. (Boeing does not agree.) He worries deeply about the Maxes still flying, saying unnerving things like, "God forbid another crash occurs." (Boeing shoots back that it has full confidence in the Max's safety; the planes carry 700,000 passengers a day.)

Pierson and Boeing have long settled into a David-and-Goliath antagonism, perhaps intensified by having once known—even respected—each other up close. This morning in DC, however, Pierson is convinced he *finally* has the receipts to make an overwhelming case. He'd explain more, but at the moment, it's time to move. Pierson walks out the hotel door. He strides south toward the brutalist behemoth at 935 Pennsylvania Avenue: home of the FBI.

TWO

"I'm mentally prepared for rejection," he explains en route. This is not Pierson's typical mode, as a former Naval commanding officer, a Division-1 football player, son of a cop—an American big guy with Big Guy Confidence. But he hasn't had time to set up meetings. No gravitas gained by his tourist's polo shirt and jeans in a citadel of suits. His wife, Michelle, advises him to take off his baseball cap before opening the FBI's door. He pushes it open, all of us curious how the first stop will go.

In a word: disastrously.

Pierson walks up to the counter, sets his backpack on the ledge, and as he unzips it, says to the man behind the window, a tad jumpily, "If you don't mind, I'm just gonna take out the package. It's just an envelope for FBI director Wray. Is that OK?" The man calls a 270-pound tank from security. That man makes Pierson step outside and stand back several feet. The guard takes a tactical position—back to the wall, sight lines on Pierson, Michelle, and me. He tells Pierson that the FBI can't accept evidence this way.

Once the guard leaves, Pierson grouses, "I hate being put in that box of 'He's a potential loo-loo, crazy person.'"

The next drop, at the Department of Justice a half block away, doesn't go much better. A pair of wary guards transmit a *Suuuuure you know someone* skepticism while Pierson leaves a voicemail for the fraud division head on the Boeing case. One guard lingers by the door of our Uber until it pulls away.

Zero for two.

Pierson dials more numbers as we drive to the next target. Now he's getting somewhere. The FAA sends down a runner—"Mr. Pierson?"—to grab his envelope in the lobby. The mail guy at the National Transportation Safety Board promises to get the documents

A 737 Max in production at Boeing's factory in Renton, Washington.

to the chair. The fraud chief at the DOJ calls *back* and directs Pierson to another building.

In the lobby of the Securities and Exchange Commission, Pierson gets back on the phone, dialing a glut of numbers off the website. Pretty soon its chief counsel appears and takes the envelope. A Department of Transportation contact meets Pierson at a Starbucks for the drop, tense that a journalist is in tow. Two staffers from the Senate Permanent Subcommittee on Investigations, which has been grilling Boeing since the Alaska Airlines door blowout, accept the envelope in the bowels of an office building. Traipsing down the hallway, Pierson whispers, "I think the committee staff think I'm kind of obsessive."

What's not clear from this one-man mission is that Pierson sits at a hub of Boeing dissent. He's become a trusted counsel to the families of Max crash victims. His podcast, *Warning Bells*, is now a way station for the company's critics, almost all of whom want Boeing to get back to its old self. He has formalized a group of experts into a foundation that parses technical reports and blasts out pugnacious letters to US agencies. Other Boeing whistleblowers have reached out in search of a sympathetic ear. Some leak new information, like these very documents he is dropping around DC.

Thousands of feet above, across the globe, more than 1,700 Max airplanes crisscross the sky, and Pierson constantly reads the malfunction reports still coming in from airlines. He worries it all bodes badly.

In aerospace circles, Pierson's name gets big reactions. Some praise him as a safety warrior, a noble soul. Others think his painting of Maxes as flying coffins is inaccurate and unhelpful, that Pierson is a hammer in search of the next Boeing nail. ("It seems he's built a little empire out of it," says one critic.) But there's a reason his supporters keep listening: They find him genuine and knowledgeable—and bad events keep happening.

"When that door blew off," a Max pilot says, "I remember talking to Ed and saying, brother, you've been affirmed."

THREE

To understand why Pierson will not let the crashes go, it helps to know he grew up with a father who solved homicides. In the 1970s, Pierson wanted to be a police officer, too. He idolized his dad, Ray, a towering DC detective who gave him plenty of *like-hell-you-will* lectures. During Pierson's senior year, Ray rustled him out of bed and drove him to the Naval Academy to take in the green grass and white uniforms and meet with a football coach. Pierson left his cop dreams behind. At the academy, he played defensive end and met Michelle, a spunky Navy brat. After graduation, Pierson became a naval flight officer (call sign "Fast Eddie") aboard a Lockheed Martin P-3C Orion.

Pierson and Michelle married in 1987. In an early assignment, Pierson worked at the Pentagon and later helped coordinate military operations with the State Department. The couple shorthanded their incoming kids with Seussian nicknames: Things 1, 2, 3, 4, and 5. In the '90s, the family settled in Washington state, near Michelle's parents. Pierson joined the reserves, rising to lead a more than 300-person squadron that ran reconnaissance missions off the West Coast. His dad showed up at the change-of-command ceremony, proud that his son's name, per tradition, adorned the side of a plane.

Pierson started his career at Boeing in 2008, with jobs in fleet services and the test flight division. In 2015 he moved to the company's factory in Renton, on the lip of glittering Lake Washington, where he oversaw managers of engineers working on assembly-line efficiencies and also managers of "shipside" teams that coordinated fixes on the floor. Pierson took in the humongous line of 737s, the company's longtime workhorse—"a kid in a candy store," as he put it, awed by the manufacturing ballet. At the time, Boeing was pushing out the first test planes for the 737 Max.

The Max was infamously rushed from the start. Boeing's archrival, Airbus, had released a more fuel-efficient model, a huge draw for airlines. Boeing executives decided they'd quickly modify the 737 to compete. The new version became the fastest-selling plane in Boeing history. By 2017, as Max production ramped up, executives called on the Renton 737 program to push out an unprecedented 47 planes a month.

> "Frankly right now all my internal warning bells are going off. And for the first time in my life, I'm sorry to say that I'M HESITANT ABOUT PUTTING MY FAMILY ON A BOEING AIRPLANE.

By the end of the year, Pierson says, he was troubleshooting the aviation version of the *I Love Lucy* chocolate factory speedup: "It was chaos." Suppliers were behind on delivering parts. On the factory floor, Pierson saw workers installing components out of the prescribed sequence. People were required to do overtime, weekend after weekend. Old-timers told Pierson this was the worst they'd ever seen: Pierson recalls that workers were logging over 30 percent more quality issues.

As 2017 turned into 2018, Pierson trudged home each night after work to his quiet house deep in Boeing country, bone tired and increasingly worried. *Where was the tipping point from a chaotic factory to a dangerous one?* (Boeing says its factory conditions did

not affect airplane safety.)

Miserable and fearful, Pierson was thinking about retiring early and finding another job. He was in his mid-fifties. In June 2018, the factory was again told to speed up production, to 52 planes a month. Pierson lost his parking spot: Unfinished 737s, including Maxes, filled the factory lots as they waited for parts to come in. One Saturday, headed in for weekend work, Pierson blazed up I-405 at dawn. He missed his exit. *If I'm making blunders, how mistake-prone are the rushed workers grinding out months of physical labor?*

Pierson decided to go straight to the top. He typed out an email to the head of the 737 program, Scott Campbell. He wrote about the 38 planes waiting for parts and the exhausted employees: "Frankly right now all my internal warning bells are going off. And for the first time in my life, I'm sorry to say that I'm hesitant about putting my family on a Boeing airplane." He asked Campbell to shut down the production line until they could finish the planes parked outside. (Campbell, since retired, didn't respond to a request for comment.)

Campbell thanked him for his "great insight," adding that he would remind leadership that safety came first. The next month, Pierson, having seen factory statistics only get worse, walked into Campbell's office and again raised his concerns—the parts delays, a lack of quality inspectors, electronic and electrical tests failing. He again demanded the line be shut down. "We can't do that," he remembers Campbell saying. "I can't do that." Pierson responded that in the military, they'd shut down missions for lesser safety concerns. He says Campbell replied: "The military isn't a profit-making organization."

A couple weeks later, in August, Pierson retired ("abandoning the *Titanic*," he'd later say). Dozens of employees penned you'll-be-missed messages around photos of Boeing planes. One coworker with whom he'd traded worries added, "I am sure that you will be happy to leave this place."

Pierson figured he'd take a break before looking for a job. In the meantime, he poured himself into his side gig on Bainbridge Island, helping coach teen football with his own high-school teammate. Yet day after day, watching defensive drills, he couldn't stop worrying about the Max factory.

In late October 2018, he sat in his family room, studying football videos on his laptop, when he overheard a headline on TV. *Plane crash in the Java Sea. Lion Air. All 189 passengers killed. A two-month-old Boeing 737 Max 8.*

Michelle walked in. Stunned, Pierson said, "Honey, this is a brand-new plane. I was there when it was being made."

FOUR

Pierson started searching, fervidly, for news. Images of gnarled plane parts being fished from the sea made him sick. One image, of shoes lined up on a pier, yanked him from sleep: tiny red slippers with Velcro—children's shoes. His horror and anger built for days, then weeks. The news reports highlighted Lion Air's checkered safety record and potential problems with the plane itself. The company assured the world the Max was "as safe as any airplane that has ever flown the skies."

A month after the crash, the Indo-

nesian government released a preliminary report. Pierson printed out all 78 pages and devoured it on a flight (not a Max; he vows to never fly in one). The report included a Boeing bulletin to pilots about how a faulty sensor could trigger the plane's nose to repeatedly point down. No mention of the factory. Pierson felt he had to draw the investigators' attention to the assembly line.

That December, "naive enough to think they'd want to talk to me," he mailed a letter directly to Boeing CEO Dennis Muilenburg; he wanted to get in touch with the company's lead on the Indonesian investigation. He wrote: "Admittedly the information I need to share isn't favorable to Boeing, but I believe it is very important nonetheless."

The response came from the company's lawyers: first, phone calls asking about his concerns, and then in February, an email saying senior leaders had looked through them: "We have seen nothing ... that would suggest the existence of embedded quality or safety issues." Pierson replied bluntly, "I don't think this is a sufficient response." He wrote to the board of directors. He didn't want to "wake up one morning and hear about another tragedy and have personal regrets." That was February 19, 2019.

Before dawn on March 10, Michelle was startled awake, as Pierson, who had been thumbing through news on his phone, roared, "FUUUCK!!"

Ethiopian Airlines' ET302, on a four-month old Max, had plowed into a field outside Addis Ababa at 575 miles per hour. All 157 people aboard were killed. Within days, regulators around the globe grounded all Maxes—more than 350 of them. The US House Transportation Committee launched an investigation and urged Boeing employees to contact them with information.

Pierson typed out a message.

Until then, his agitating had taken place behind corporate doors. Now he hired an attorney. Over the next weeks and months, Pierson told a slew of federal agencies what he'd witnessed at the Boeing factory. Investigators from the House Transportation Committee reached out to talk, and by fall, they wanted Pierson to testify at a public hearing. Doug Pasternak, the lead investigator, saw Pierson as "the poster boy of whistleblowers"—unimpeachable résumé, meticulous documentation, a keen sense of moral duty. "Ed was the only one," Pasternak says, "where we had evidence that he was screaming from the rooftops to senior management."

Pierson understood his presence on Capitol Hill would make an impact—and bring a new level of scrutiny. So he and Michelle phoned each of their kids, who were now in their twenties and thirties: This would be Dad's face on national TV, speaking damningly about a powerful company. What did they think? They all replied, *Do it.*

Quickly, journalists caught wind of the Boeing insider about to go public. Pierson spoke with The New York Times. A documentary crew flew with him and Michelle to DC, filming Pierson's brooding and Michelle's face masked in worry. In DC, he sat before NBC's crew: "I'm mad at myself, because I felt maybe I could have done more."

In the hearing room at the Rayburn congressional building on Capitol Hill, Pierson sat behind a microphone. His voice had the haunted tone of a eulogy. He laid out his history of raising concerns while on the job and after the first crash, and asked for the factory conditions to be investigated. In the gallery, among the parents, sons, and daughters of crash victims, a man from Toronto named Chris Moore held a photo of his 24-year-old daughter, Danielle, who had been killed in the Ethiopian crash. Listening, Moore thought *here*, finally, a guy with the moxie to raise safety issues *before* the disasters—and it was the first he'd heard about Renton as a potential factor. After the hearing, the mother of another victim hugged Pierson, telling him, "That took courage."

"Why is she hugging *me*?" Pierson wondered. "I worked for the company that built these damn planes."

The Piersons flew back home. "I hoped," Michelle says, "that now that he'd given his stuff to Congress, it would be done. He could be free of it."

Instead, Pierson burrowed in.

FIVE

In the news, Boeing's spokesperson gave Pierson a pat on the head: He "did the right thing" and the company "took appropriate steps" to address his concerns. But the spokesperson said that Pierson's allegations were "completely unfounded" and that no investigator had linked the crashes to factory issues.

At this point, it was widely reported that the cause of the crashes was a series of events triggered by faulty angle-of-attack sensors. Mounted on the plane's nose, these sensors calculated the angle of the wings to the oncoming air. The devices were made by another aerospace company, and Boeing workers normally installed and tested them in Renton. Each Max had two sensors, but Boeing designed the system so that only one sensor fed data into the flight control computer at a time. In the minutes leading up to both crashes, that lone sensor had failed; it gave the pilots a false reading that the plane was angling sharply upward. That reading repeatedly triggered a software system called the Maneuvering Characteristics Augmentation System. To prevent the plane from stalling if it were actually angled dangerously upward, MCAS automatically pitched the plane's nose down. Pilots on both flights fought the automatic system in a chaotic tug of war before finally losing control.

"From day one in aerospace engineering undergraduate classes," says MIT aeronautics lecturer Javier de Luis, who lost his sister Graziella in the Ethiopian crash, "we tell kids you cannot have single points of failures in systems that can take down an airplane." Having one sensor transmitting data into MCAS created a single point of failure.

Boeing also hadn't been fully transparent about MCAS with regulators and didn't tell pilots of its existence. The FAA requires that critical software like MCAS undergo more vetting and that airlines potentially do more pilot training. Journalists and investigators roundly dissected how Boeing downplayed MCAS to the FAA and marketed the Max as needing minimal training for pilots of other 737 models.

Pierson fixated on why the sensors, lost in the wreckage, had failed. The day before the Lion Air flight, mechanics in Indonesia had replaced one faulty angle-of-attack sensor with another, refurbished one. In the final crash report from Indonesia, the authors blamed the refurbished sensor's failure on a Florida repair shop. The shop, they wrote, had miscalibrated it by 21 degrees. The FAA revoked the Florida shop's certification to repair airplane parts.

As for the Ethiopian plane, the National Transportation Safety Board (NTSB) and its French counterparts, who'd gotten involved in the investigation, attributed the sensor's malfunction to it getting physically walloped. Most likely by a bird. (In fact, these sensors have failed after being frozen, installed incorrectly, hit by lightning, or, yes, whammed by birds.) The Ethiopian investigators, however, were skeptical from the beginning. They would eventually report that they hadn't found any bird carcasses. (American investigators argued that the Ethiopians took more than a week to check for bird evidence.)

Pierson accepted that the sensors had failed. It was the "why?" that bothered him. He couldn't shake his sense that there was more to the story.

SIX

In early 2020, as Covid clobbered the airlines, the media's glare on his congressional testimony dimmed and Pierson bought an RV. His hope was to road-trip to Texas to see his grandkids. He never made it, but he did park the RV at his in-laws' farm nearby and practically moved in. For up to 12 hours a day, Pierson sat at the RV's kitchenette table, combing through crash reports, news stories, electrical engineering tomes. He etched notes and Ohm's law into a giant artist sketchbook; he plugged any mention of the Max's sensors into an Excel spreadsheet.

Surveying his mess of papers, Pierson's thoughts swerved to his dad back in the day, always mulling his latest whodunit. His father had died more than a decade prior. But there in the RV, in the silence of the Covid summer, Pierson felt close to him. He thought about something his dad always said: People lie; follow the evidence.

By September, Pierson was stewing. The FAA was getting close to putting the Max back in the sky. Boeing was settling lawsuits with the Lion Air victims' families, and it had changed its problematic software. The head of the FAA touted the agency's recertification work and piloted a test flight himself.

That month, the House Transportation Committee finally released its 245-page Max report. It outlined how Boeing obscured information about the software from pilots; it criticized the FAA for failing to act after the first crash. Pierson's name appeared 132 times—an "incredibly important" source, the committee wrote, who painted "a deeply troubling picture of Boeing's production first, safety second, culture among Boeing's senior leadership."

In his RV, Pierson plowed ahead. He reread the 322-page final report from Indonesia and an interim report from Ethiopia. He focused on something that was nearly absent from the public narrative: The new planes had maintenance problems, including likely electrical issues. The Lion Air plane had a tripped circuit breaker and a flurry of fault messages, one of which prescribed wiring checks. On the Ethiopian plane, the auxiliary power unit had malfunctioned, and the captain's computer outlet never worked. The altimeter and vertical speed indicator were giving erratic readings. Dramatically, while on autopilot in the months before the crash, the plane rolled to the side unprompted.

> "Wow, it's almost like someone's been **SCREAMING TO DEAF EARS FOR FIVE YEARS** about production issues.

"You sons of bitches," Pierson remembers thinking, "this information should have been pursued."

As he scrutinized the Indonesian report, nuances emerged. Investigators first suggested the *possibility* that the Florida shop miscalibrated the sensor; they *could* have inadvertently introduced a 21-degree error. As Pierson read, he noticed that in other parts of the report, the idea hardened into firm fact.

Then something deep in the report grabbed his interest. Several months after the crash, investigators from the NTSB, Boeing, the FAA, and the sensor manufacturer examined the *old* sensor from the Lion Air plane, the one that had been replaced the day before the crash. *That* sensor had signs of electrical damage, including a broken wire and an intermittent open circuit.

Pierson got in touch with an aerospace engineering professor named Daniel Ossmann at the Munich University of Applied Sciences, an expert on detecting faults in airplane sensors, and peppered him with technical questions on a series of Zooms.

Pierson settled on a theory. Maybe the sensor wasn't hit by a bird in Ethiopia, and the Lion Air Max's sensor wasn't a lemon. Maybe a malfunction in the plane's electrical infrastructure *made* them lemons, such that they gave bad readings. Or the reverse: A damaged sensor somehow hurt the electrical system.

He wrote a report, called "Still Not Fixed," and sent it off to Ossmann and other aviation experts for feedback. In January 2021, he posted it on a vintage-looking website he'd hastily built on GoDaddy. With meticulous footnotes and links, Pierson outlined new lines of inquiry about the electrical system, the miles of wires that form the plane's nerve system. He noted that he saw electrical defects getting detected and reworked at the Max factory.

Pierson was taking a firm step away

The factory "chaos" led Pierson to retire early. Colleagues wrote well wishes on a photo of a plane.

from government findings, and, at that point, doing it alone: US authorities stuck to miscalibration and a bird strike to explain the sensor failures. Boeing, citing those authorities, told WIRED, "As we said in 2019, Mr. Pierson's theory is wrong." Pierson argues that he "never claimed that I know everything about what could have happened electrically." What he wants is a more detailed investigation.

Pierson sent his treatise to journalists who had covered the Max debacle, hoping for another blast of press.

The BBC did a story, quoting a plane safety expert who cautioned that interpreting the crash reports doesn't constitute a new investigation. But most reporters had moved on. Earlier that month, the Department of Justice charged Boeing for conspiracy to defraud the FAA about its MCAS software. At the same time, the DOJ announced that the case was being settled: Boeing agreed to ramped-up federal oversight, a $243.6 million penalty (roughly the sticker price of a wide-body Boeing airplane around that time), and $1.77 billion paid to airlines for lost revenue. Boeing would also pay $500 million to victims' families.

The government wrote, "Misconduct was neither pervasive across the organization, nor undertaken by a large number of employees, nor facilitated by senior management." Boeing touted its Max changes; global regulators deemed it safe to fly. And Pierson went back to his RV, restless.

Pierson had spent his life contributing to institutions. At his most senior post in the Navy, he'd coordinated operations for 350,000 military personnel. Now he'd even quit coaching football to focus on Boeing. He and Michelle lived off of his military and Boeing retirement pay and their investments. He'd sold all but a few Boeing shares, keeping those to spy on shareholder meetings. He had become a lone wolf—or worse, a gadfly. The only people who listened to him were his kids, Navy and football buddies, and, always, Michelle.

Pierson sent his report to a Senate committee working on aviation reforms, which interviewed him and attached his treatise to its aviation safety whistleblowers report. In back-channel conversations, some people found Pierson's ideas plausible; others dismissed them. "Of course we listened to him," says one person close to the Senate investigation, "but it was hard to validate that production was the cause of the Max crashes." Over at the NTSB, leaders would send him three letters saying that the evidence didn't point to the factories. "He just kept on going," says a former NTSB investigator. "He's a bit of a one-trick pony."

Pierson wasn't only upset about losing momentum. As his phone went quiet and people again boarded Maxes around the world, Pierson felt lonely.

At the end of the year, he connected with a crash victim's mother. She was Ralph Nader's niece, Nadia Milleron, and he'd talked to her a few times before. Ever since her daughter died in the Ethiopian crash, she had lobbied relentlessly for regulatory changes. Milleron frequently Zoomed with a few families who were fighting for aviation safety. On the phone, she told Pierson they could use some technical help. So in December 2021, Pierson nervously signed in to their Zoom, bracing for a potentially cold reaction ("These people probably hate me") and an unfathomable well of grief.

SEVEN

Nadia Milleron's daughter, Samya Stumo, had been 24, bright, altruistic, with "a spreadsheet for every occasion." She lived in DC and was traveling to East Africa to set up an office for the health-focused nonprofit where she worked, telling Milleron over the phone before she left, "It's just two weeks, Mom." After the crash, Milleron experienced cascading effects of trauma. Sometimes she felt like she couldn't feel her feet while walking—as if "paddling through the air" of her Massachusetts farmhouse. She'd set a timer to measure how long she was scrubbing dishes. She'd count aloud while making tea, while driving, to keep herself tethered to time. One night, she looked through the catalog of items collected from the crash debris, and she saw Samya's blue skirt. When it arrived, it smelled like jet fuel.

Milleron found some purpose in agitating for reforms, especially with other families. Another of those agitators was Chris Moore, a retiree in Canada. Signing in to Zoom, Moore recognized Pierson from that congressional hearing two years prior, in which he'd listened from the gallery while holding a photo of his daughter, Danielle. She had been a climate activist and taught kids to code. She used to bike through the Manitoban winter and tweet out selfies of the icicles on her eyelashes. Hours before the crash, she posted: *"I'm so excited to share that I've been selected to attend and am currently en route to the @UNEnvironment Assembly in Nairobi, Kenya."* Every evening at 6 pm, Moore and his wife, Clariss, light a candle while they cook dinner. He still uses the email address Danielle made for him in seventh grade—"dadthedude." He thought, often, that if he couldn't wish Danielle off the plane, he wished himself onto it, to hold her hand through the horror. "Why *my* daughter? Why did it have to be *her*?"

Pierson—dad of five, granddad of three so far—took in the handful of faces on the Zoom grid. He expressed his sorrow for their losses and said he had information to share. The families welcomed him. From that point on, Pierson answered or researched their questions and joined their regular calls. Could Pierson sometimes go a little long on his passionate jags? Sure, "but those are style things," Milleron says. "I see Ed as a key. As a clue." Moore calls Pierson "your typical GI Joe," a trustworthy wingman with a reassuring refusal to fear a corporate titan. And Pierson, "he probably feels that his sails are full from us, because we believe in him. We agree with him." (For Moore, that includes supporting Pierson's electrical theory.) In his loneliest year, Pierson had found another group that had very much not moved on.

As Pierson plunged forward in his research, he only grew more anxious. Earlier in 2021, Boeing had temporarily grounded more than 60 Maxes for a new electrical problem. At one point, Pierson had called up a manager who had reported to him in Renton. The manager told Pierson that both the Ethiopian and Indonesian planes showed hydraulic, electronic, and electrical problems during manufacturing, and that employees had logged requests for assistance with those issues into Shipside Action Tracker records. The problems were supposedly fixed before the planes went into service. Pierson asked to see them—"I told Michelle, we'd mortgage the house to get those SAT records"—but the manager said he didn't save copies.

Pierson was buoyed by a development at the end of 2022: The Ethiopians released their final crash report. Pierson had reached out to them amid their investigation and sat for more than a dozen interviews. At one point, he even sent them questions to ask the NTSB. Now the final report was almost exactly aligned with his theory. Citing the plane's electrical abnormalities in the weeks before the crash, the report claimed that electrical defects from the plane's manufacturing caused the sensor to fail—not only in the crash they were investigating but in the Lion Air crash too.

The NTSB and the French safety agency quickly shot the report down, a rare disagreement between international aviation authorities. They repeated the bird-strike theory; the French even claimed an object hitting the sensor was "the only possible scenario." The NTSB criticized the Ethiopians for not examining the pilots' performance. Some cautioned, again, that the real problem was that Boeing designed a system with a single point of failure that—whatever its cause—could take down planes.

Pierson, of course, remained firmly on the side of the Ethiopians, but his focus had widened to the Maxes in the air, the ones built alongside the doomed planes and since. And he was no longer alone.

Another whistleblower, Joe Jacobsen, had joined the family calls. Also in his sixties, Jacobsen had retired from the FAA in 2021, disillusioned with what he says was the agency's kowtowing to Boeing. As a safety engineer, he'd been in an FAA meeting where Boeing engineers presented the idea of the bird strike, and he'd immediately had doubts. (Among the reasons for his skepticism: In the plane's recovered cockpit recorder, there was no thump or verbal acknowledgment from the pilots that they'd hit a bird.)

Jacobsen and Pierson started poring over information together—including a 343-page Boeing bulletin from 2020 that prescribed changes to a specific Max electrical system to comply with regulations. They kept a list of Boeing's petitions to the FAA for exemptions from design safety standards. They eyed new problems on Maxes that were being logged in two arcane federal databases: Switches that didn't work. Airspeed indicators that didn't agree. Emergency landings that never made the news.

Tired of courting the media for attention, Pierson started his own podcast, *Warning Bells*. On a couple episodes, Pierson mentioned that a single airline, Alaska, was reporting more problems with the Max than others. All Maxes come off the same assembly line, so he didn't think it made sense for Alaska to have more issues. He figured that the company was simply more diligently reporting them. In the spring of 2023, he FedEx-ed a letter to an Alaska Airlines PO box, addressed to the CEO, in which he recommended that Alaska ground the planes. (The airline says it never got the letter.) After Pierson put out a press release on the issue that fall, the reports from Alaska sharply fell off—from more than 90 a month to a trickle. Alaska says that it changed its reporting "to the industry standard," logging certain types of incidents to the FAA internally rather than in the public logs Pierson looked at.

By now, Pierson had formed a tight group to compare notes with, including some Navy buddies and a pilot's union spokesman who flew Maxes and had publicly criticized Boeing. Plus Jacobsen and another retired FAA whistleblower, Mike Dostert, who had warned his managers about the risks of "loving Boeing to death."

That summer Pierson pitched them all an idea: How about joining up as an independent watchdog of Maxes and other planes? An expert "fighting force," Pierson says. The group set up a nonprofit. A friend—another football coach—put up a sharp website: the Foundation for Aviation Safety. They issued a first press release about faulty electrical bonding on Maxes. It didn't exactly grab headlines.

On the other hand, a door flying off a plane does.

Pierson and Jacobsen were sitting on the bleachers in a school gym in January 2024 when Pierson's phone buzzed with texts and news. A door plug had ripped out of a new Max mid-flight, high above the city of Portland; passengers strapped on oxygen masks. Alaska Airlines.

Jacobsen turned to Pierson to say, "I think we're about to get busy."

EIGHT

Suddenly, many of the moves Pierson had wanted started to happen. The FAA grounded Max 9's—the larger model, like the one that had lost a door—and the NTSB marched into the Renton factory.

Within a month, they announced that Boeing workers hadn't installed four bolts that hold the door panel. Inspections of other grounded Max 9's turned up loose bolts. The FAA and US senators started talking about a "systemic issue" with Boeing's manufacturing. The FAA gave the company 90 days to fix its quality control.

Pierson's daughter on Facebook: "Wow, it's almost like someone has been screaming to deaf ears for 5 YEARS about production issues?" To me, Michelle snuffed any notion that her husband would gloat: "Do you want to be right about that? He was just grateful that the whole plane didn't go down." Even one of Pierson's critics, a former NTSB investigator, conceded that on this one, "Yep, that's a manufacturing quality-control, safety-culture problem, and Ed Pierson got that right."

Pierson was once again, for the first time in years, flooded with media requests. Four days after the blowout, he spotted an email from John Barnett. He recognized the name: Barnett was a former Boeing quality manager who

If a corporate and regulatory culture of yes-men and -women led to the deaths of 346 people, then Pierson will happily be the nope man, **AWARDING NO BENEFIT OF THE DOUBT.**

had spoken up inside the company and to regulators about alleged shortcuts at Boeing's 787 Dreamliner factory in South Carolina. He'd been locked in a years-long legal battle with Boeing over alleged retaliation. He now told Pierson he'd seen him on the news and was interested in working with his foundation. Pierson liked the idea of Barnett joining. Figuring they'd probably talk for a long time once things calmed down, he moved the message into his "Follow Up" folder and returned to his sprint.

In mid-March, Pierson opened an email: *Had he seen the news?* John Barnett was dead. He had taken his own life, sitting in the driver's seat of his truck amid a multiday deposition in his case. (Barnett's family is continuing on with the retaliation case.) Reeling, Pierson immediately dug up Barnett's email to see if he had missed a cue. The tone was gracious and professional. Barnett wanted "to advocate for safety" and was asking "if there is anything I can help with." Pierson swirled in regret. "If I'd called him back sooner, would it have made a difference?"

Later that spring, Barnett's death rippled through a Senate committee hearing on Boeing. Pierson raised his hand to take the oath. His 2019 testimony had been laced with heartbreak, but now he was punchier, angrier, as he let loose his controversial take: "The manufacturing conditions that led to the two 737 Max disasters also led to the Alaska blowout accident, and these conditions continue."

Boeing's internal safety line was exploding. So was the FAA's complaint line. With Pierson and other foundation folks all over the news, emails sailed in from former and current Boeing employees who wanted to commiserate. Thinking of Barnett, Pierson tried to get back to them quickly. Some Boeing dissidents invited Pierson to what they called the "Ed Head Picnic," named in his honor. While yet another documentary filmmaker captured the gathering, the organizers knighted him "Whistleblower One" and hung a whistle around his neck.

NINE

It was late at night when Pierson saw the email from a Boeing employee. He clicked open one hell of an attachment: an internal Boeing record about the doomed Ethiopian Airlines plane.

Pierson burrowed under his blanket to shield Michelle, sleeping, from the glare of his phone. He zoomed in on a Boeing technical support log. It was from December 2018, not long after the Lion Air crash. It showed that the Ethiopian Airlines Max, which had been in service only a few weeks, had spontaneously rolled to the right while on autopilot. Pierson remembered that roll from the crash report. In this record, Ethiopian Airlines had asked Boeing to explain the cause and offer a fix.

What was new to Pierson was the Boeing engineer's reply: Boeing had come to "suspect" an "intermittent fault" with some wiring to a flight control computer where MCAS resided and a component that receives angle-of-attack information. Boeing advised the airline to inspect specific wires for shorting or an open circuit. The plane crashed three months later.

Pierson was aghast. Concerns about the wiring, right there in words. He emailed the document to the Ethiopian crash investigators—the ones who'd blamed electrical defects, and who US and French authorities had dismissed. An investigator replied that he'd never seen this record. "I said, 'How important would it have been to your investigation?'" Pierson recalls. "He said it was vital." (Boeing says it shared information from this document with investigators, but didn't answer a question on whether the company shared the document itself.)

A few months later, another email from the same Boeing insider: This one contained the Ethiopian plane's manufacturing troubleshooting records that Pierson had chased years earlier—the Shipside Action Tracker records—as the plane moved down the Renton factory line in the fall of 2018, right after Pierson had retired. Reading through the entries, it seemed to Pierson that employees were confused about what electrical parts had been installed. They seemed to be talking about mislabeled components and a botched wire bundle job, plus pressure to get overdue parts from Boeing's electrical device manufacturing shop in Everett, Washington. That shop had its own quality problems at the time. Responding to whistleblower tips, the FAA interviewed the shop's staff and confirmed that employees had done unauthorized work on certain parts after quality inspectors had signed off, in violation of FAA rules. (The complaint doesn't state which airplanes those parts were destined for.)

To Pierson, it all bolstered his contention that electrical problems could have caused the Ethiopian plane's sensor failure—or, at least, that the notion should have been better examined.

Gathering the documents in late July, Pierson was summoning his football metaphors, saying he might make a "touchdown in the final seconds." A deadline loomed: The US government had offered Boeing a new plea agreement after the door blowout, and the victims' families—who blasted it as a "sweetheart" deal—were due to file their opposition. Pierson thought the families might want to use his new information in their arguments.

Pierson and his wife, Michelle, in the woods near their home outside Seattle.

As it happened, the day before the families' legal deadline, Pierson was slated to speak at the National Whistleblower Day in Washington, DC. So he and Michelle got an earlier flight. He didn't want to be honored for past whistleblowing as much as do the actual thing: drop his documents straight into the hands of power.

TEN

Part of the motive for this grand dash was Pierson's old-school flair: Emails get ignored. He likes the implicit pressure of showing up in person. (Pierson's motto is "Red, white, and blue, this is the right thing to do," said a former congressional staffer who has talked with him for years.) Pierson didn't have time to set up meetings, so he waltzed door to door, from agency to agency, and worked his phone in the passenger seat of Ubers as he crisscrossed DC. Finally, back at his hotel that afternoon, he sat down and emailed the documents to the families and their attorney. After Milleron read them, "I was a complete mess," she told me. "They could have taken that plane out of service and Samya would be alive."

The next morning, at the National Whistleblower Day celebration, the Piersons filed into the grand marble chamber that once hosted the Watergate hearings. Michelle fastened a "Celebrate Whistleblowers!" pin to her purse.

More than 100 people filled the room: Whistleblowers from Exxon Mobil, MIT, a glut of federal agencies. Enron whistleblower Sherron Watkins heartily shook Pierson's hand. Big Tobacco whistleblower Jeffrey Wigand, now in his eighties, joked to the crowd that he was actually Russell Crowe, who played him in *The Insider*. As a C-SPAN camera rolled, Pierson, now in a suit, held up his documents at the podium and announced his drop to the nation's power brokers. Audience members jumped to their feet, clapping—a whiplash in public regard, given that 24 hours earlier he'd gotten bounced at the FBI.

This past October, Pierson and Jacobsen flew to Fort Worth, Texas, for a hearing about the proposed "sweetheart" plea deal for Boeing. The terms on offer: plead guilty to conspiring to defraud the US, pay another $243.6 million fine, and accept three years of probation. Nadia Milleron and her husband flew down from Massachusetts, joining other victims' relatives from Ireland, France, and California. Chris Moore and his wife, Clariss, drove all the way from Toronto. (He boycotts flying.) Sitting down, the families spotted the foundation guys in the gallery and waved. Pierson and Jacobsen had shared their new documents directly with the judge, alleging they showed "Boeing's continuing concealment of critical public safety information." (The judge would reject the deal, and at press time, Boeing and the government were still talking about the resolution.)

After the October hearing, the families joined Pierson and Jacobsen at a Mexican restaurant. A boom mic from a documentary crew hovered above Pierson's head. Jacobsen pulled out a suitcase from under the table, and Pierson handed out glass awards, from their foundation, honoring the families' leadership on aviation safety. Pierson improvised a speech for each one.

Chris Moore thought, well, this was unexpected. "You don't think, oh, I can't wait to get an award someday." But at this point in the awful five-year battle that he never wanted, "shaking my fist at the clouds," as he put it, a token for the Zoom group's efforts felt nice. Moore knows that all this fact-finding and accountability-seeking serves another purpose, too: to help protect him from his bottomless grief.

Pierson still wrestles with his own grief, a wholly different kind. Could he have done more to prevent the crashes? "I don't think I'll ever—" He lets out a long exhale. "I'll ever stop feeling that way."

Listening, I thought about something Doug Pasternak, the lead investigator of the Max report, told me about his conversations with Pierson. "He was devastated. He did have a sense of, 'guilt' may not be the word, but responsibility. He just wishes there was something that could have been done to prevent these horrific accidents."

Pierson couldn't prevent the crashes, although no one I spoke to thought he could have done more. But he could become the guy hellbent on not letting another Max fall from the sky. He could hunch over every report to work out possible explanations in an RV kitchenette. He could be the fired-up guy pushing authorities to look—no really, *look*—under every last Boeing rock. If a corporate and regulatory culture of yes-men and -women led to the deaths of 346 people, then Pierson will happily be the nope man, awarding no benefit of the doubt.

The new documents, with all their promise of bringing home Pierson's contested electrical theory, ended up amounting to less than he'd hoped. The NTSB told Pierson it wouldn't hand the papers to the Max crash investigators—the cases had concluded, the board said—but he could do so himself.

Boeing wobbles in limbo, before civil and criminal courts, at the FAA, in Congress, awaiting the final doorplug report from the NTSB. Observers say 2025 will be Boeing's pivotal year: The company either turns around under its new CEO or succumbs to a doom loop. Pierson vows to keep talking.

"For me, it was always about not allowing them to shut me up," he says. Recently, the foundation received its first donations and now has a payroll. They're starting to monitor other aircraft models and are talking with a university about analyzing industry-wide data—"to be an equal-opportunity pain in the butt," Pierson says. The guy Boeing surely hoped would go away by now has, instead, institutionalized himself to stick around.

When Pierson said goodbye to me in DC, his parting words were: "Don't fly the Max." I couldn't bring myself to tell him. That's exactly what I was booked on, the 7:41 pm from Dulles to San Francisco. It was the one I could catch after the whistleblower event on Capitol Hill and still walk into my house that night. Commercial flight was supposed to be about convenience, after all, collapsing a country's span into a Tuesday night commute. At this point in aviation history, we passengers *should* be able to pick a flight on time alone.

Hurtling through the air that evening in seat 10C, I read the US House committee's Max investigation, a disruptor of illusions. Like many fliers, I'd long ago made my bargain with risk. I'd taken comfort in statistics, summoned faith in the engineers and assembly workers, the pilots, the system. I'd shunted away the knowledge—paralyzing, if you let it in—that stepping on an airplane is an extraordinary act of trust. Deep in the report, I reached the part about a senior manager at Boeing's factory in Renton, a guy named Ed Pierson, who seemingly knew what we all know when we soothe ourselves by thinking, *They wouldn't let it fly if it weren't safe*. We're all relying on someone to be the "they."

LAUREN SMILEY *is a San Francisco–based journalist. She can be reached at laurensmileywriter@protonmail.com.*

COLOPHON

Digital or IRL breakthroughs that helped get this issue out:

Finally figuring out how to sleep; working flexibly from home at all hours to the tune of trashy TV; banana soft-serve from Banan on 24th Street; binge-watching *Insecure* and *Only Murders in the Building*; getting a real-life boyfriend; "Baked Baking" from @mattbellassai; figuring out the best white-noise setting; Afghan Burrito in Berkeley, MexiHibachi in Pinole, Shan Dong in Oakland, Panchita's Pupusería & Restaurant in San Francisco, Isla Restaurant in Concord, and Hokkaido Ramen Santouka in San Jose; Golden Gate Angling & Casting Club; signing up for ChatGPT Plus; deleting TikTok and X from my phone; breaking a land-speed record in getting a 14,000-word feature story written, edited, fact-checked, arted, and published in under a month; finishing *The Conspiracy Against the Human Race*, by Thomas Ligotti; consistently reading the Sunday paper; the harder the task, the better my (otherwise maniac) dog is at it.

WIRED is a registered trademark of Advance Magazine Publishers Inc. Copyright ©2025 Condé Nast. All rights reserved. Printed in the USA. Volume 33, No. 5. WIRED (ISSN 1059–1028) is published bimonthly by Condé Nast, which is a division of Advance Magazine Publishers Inc. Editorial office: 520 Third Street, Ste. 460, San Francisco, CA 94107-1815. Principal office: Condé Nast, 1 World Trade Center, New York, NY 10007. Roger Lynch, Chief Executive Officer; Elizabeth Herbst-Brady, Chief Revenue Officer; Jackie Marks, Chief Financial Officer. Periodicals postage paid at New York, NY, and at additional mailing offices. Canada Post Publications Mail Agreement No. 40644503. Canadian Goods and Services Tax Registration No. 123242885 RT0001.

POSTMASTER: Send all UAA to CFS (see DMM 707.4.12.5); NONPOSTAL AND MILITARY FACILITIES: Send address corrections to WIRED, PO Box 37617, Boone, IA 50037-0662. For subscriptions, address changes, adjustments, or back issue inquiries: Please write to WIRED, PO Box 37617, Boone, IA 50037-0662, call (800) 769 4733, or email subscriptions@WIREDmag.com. Please give both new and old addresses as printed on most recent label. First copy of new subscription will be mailed within eight weeks after receipt of order. Address all editorial, business, and production correspondence to WIRED Magazine, 1 World Trade Center, New York, NY 10007. For permissions and reprint requests, please call (212) 630 5656 or fax requests to (212) 630 5883. Visit us online at www.WIRED.com. To subscribe to other Condé Nast magazines on the web, visit www.condenast.com. Occasionally, we make our subscriber list available to carefully screened companies that offer products and services that we believe would interest our readers. If you do not want to receive these offers and/or information, please advise us at PO Box 37617, Boone, IA 50037-0662, or call (800) 769 4733.

WIRED is not responsible for the return or loss of, or for damage or any other injury to, unsolicited manuscripts, unsolicited artwork (including, but not limited to, drawings, photographs, and transparencies), or any other unsolicited materials. Those submitting manuscripts, photographs, artwork, or other materials for consideration should not send originals, unless specifically requested to do so by WIRED in writing. Manuscripts, photographs, artwork, and other materials submitted must be accompanied by a self-addressed, stamped envelope.

THE ASSIGNMENT: IN SIX WORDS, WRITE A STORY ABOUT ENTANGLED PARTICLES.

HAUNTED DREAMS FOLLOWED OUR FIRST EMBRACE.

—@instaduncc, via Instagram

Honorable Mentions

Yikes. Spooky action. Time to split.
—@FilmMartin, via X

Are you here? Are you not?
—@jessleycegui, via Instagram

Breaking news: Sentient entangled particles divorce.
—Rami, via email

We spin in unison, galaxies apart.
—Mark Richardson, via email

Spooky out here! Spook, you there?
—Andrew Dawson, via email

Once it died, you were born.
—@bietorres, via Instagram

Tapestry of space, matter sewn together.
—@dr.karenorjuela, via Instagram

Meet me beyond the double slit.
—@javirz, via Instagram

Want to submit a six-word story for us to consider? Look for the latest prompt on Instagram, X, and WIRED.com/six-word, where you can also see how we've illustrated past favorites.

Disclaimer: All #WIREDSIXWORD submissions become the property of WIRED. Submissions will not be acknowledged or returned. Submissions and any other materials, including your name or social media handle, may be published, illustrated, edited, or otherwise used in any medium. Submissions must be original and not violate the rights of any other person or entity.

WIRED

WIRED SUMMER LAB

EXPLORE THE FUTURE OF TECHNOLOGY AND TURN YOUR IDEAS INTO INNOVATIVE SOLUTIONS

JULY 2025 | LONDON, UK

Designed for college students aged 18-21, WIRED Summer Lab is a four-week learning experience led by WIRED experts who will help you build tech industry knowledge and key skills for navigating your next steps in the field.

LEARN MORE AT
WIREDEDUCATION.COM

From your business goal to their first game-winning goal, every move matters.

When you partner with Bank of America, our financial expertise and digital tools can help you envision a bolder future for your business. **What would you like the power to do?®**

Visit bankofamerica.com/bankingforbusiness.

CLUB WORLD CUP 25 FIFA

BANK OF AMERICA

OFFICIAL BANK

FIFA

"Bank of America" is the marketing name used by certain Global Banking and Global Markets businesses of Bank of America Corporation. Lending, other commercial banking activities, and trading in certain financial instruments are performed globally by banking affiliates of Bank of America Corporation, including Bank of America, N.A., Member FDIC.
©2025 Bank of America Corporation. All rights reserved.